

SACO: Sequence-Aware Constrained Optimization Framework for Coupon Distribution in E-commerce

Li Kong^{1*}, Bingzhe Wang^{1*}, Zhou Chen², Suhan Hu¹, Yuchao Ma¹, Qi Qi^{1,4,5†}, Suoyuan Song³, Bicheng Jin³

¹Gaoling School of Artificial Intelligence, Renmin University of China, China

²School of Economics and Management, Southeast University, China

³ByteDance, China

⁴Beijing Key Laboratory of Research on Large Models and Intelligent Governance

⁵Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE

beverly_kong@163.com, {wbz2022, 2023202249, yc.ma, qi.qi}@ruc.edu.cn,

chenzhou@seu.edu.cn, {songsuoyuan, jinbicheng}@bytedance.com

Abstract

Coupon distribution is a critical marketing strategy used by online platforms to boost revenue and enhance user engagement. Regrettably, existing coupon distribution strategies fall far short of effectively leveraging the complex sequential interactions between platforms and users. This critical oversight, despite the abundance of e-commerce log data, has precipitated a performance plateau. In this paper, we focus on the scene that the platforms make sequential coupon distribution decision multiple times for various users, with each user interacting with the platform repeatedly. Based on this scenario, we propose a novel marketing framework, named Sequence-Aware Constrained Optimization (SACO) framework, to directly devise coupon distribution policy for long-term revenue boosting. SACO framework enables optimized online decision-making in a variety of real-world marketing scenarios. It achieves this by seamlessly integrating three key characteristics, general scenarios, sequential modeling with more comprehensive historical data, and efficient iterative updates within a unified framework. Furthermore, empirical results on real-world industrial dataset, alongside public and synthetic datasets demonstrate the superiority of our framework.

Introduction

In the highly competitive online ecosystem, coupon distribution has emerged as a central marketing strategy for online platforms looking to amplify user engagement and optimize revenue streams. Coupons take many forms (as shown in Figure 1), from traditional discounts to service upgrades and cash rewards, all designed to incentivize consumption across diverse sectors. For instance, Booking’s promotional activities can enhance user satisfaction (Goldenberg et al. 2020; Albert and Goldenberg 2022); Airbnb employs dynamic pricing (Ye et al. 2018); Uber’s promotions encourage users to adopt new products (Du, Lee, and Ghaffarizadeh 2019); Taobao’s discounts increase user activity (Zhao et al.



Figure 1: Complex and Meaningful Coupon Distribution.

2019; Zhang et al. 2021); Kuaishou offers cash rewards to stimulate user retention (Ai et al. 2022); and Meituan utilizes spend&save discounts to incentivize purchases (Zhou et al. 2023; Huang et al. 2024; Zhou et al. 2024).

However, these marketing campaigns come with significant costs. Typically, dedicated budgets are set to ensure campaign sustainability. Effective coupon distribution within these finite budgets is therefore a critical challenge, frequently formalized as a constrained optimization problem. Unlike traditional resource allocation, coupon distribution under budget constraints faces two significant challenges. First, user responses are variable and the counterfactual outcomes are fundamentally unobservable. Second, user responses are sequential, which means past coupon decisions affect users’ subsequent responses, making counterfactual prediction challenging. As a result, it is impossible to obtain a direct solution.

While prior research has explored counterfactual policy optimization from logged bandit feedback (Swaminathan and Joachims 2015b,a; Joachims, Swaminathan, and De Rijke 2018), the bandit approach relies on online exploration, rendering it unsuitable for coupon distribution in e-commerce scenarios where decisions, once made, are difficult to reverse. In contrast, more recent studies acknowl-

*These authors contributed equally.

†Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

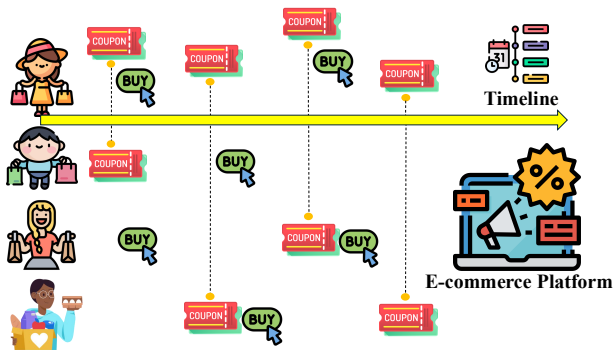


Figure 2: Single-Round Scenarios in Existing Research.

edge the importance of pre-training models with logged data before deployment, representing a significant improvement. However, they typically focus on scenarios where the platform makes a single-round coupon allocation decision for each user interaction according to the current state. The prevailing two-stage approach dominates coupon allocation research (Hansotia and Rukstales 2002; Zhao et al. 2019; Goldenberg et al. 2020), comprising: (1) machine learning (ML) prediction of user response to different coupons, and (2) optimization research (OR) for revenue maximization. However, decoupling the process into two stages leads to error accumulation, and the two-stage goals cannot be fully aligned. Recent end-to-end frameworks, including decision-focused learning (DFL) (Zhou et al. 2023, 2024) and reinforcement learning (RL) (Xiao et al. 2019; Zhang et al. 2021), attempt to address this limitation. However, these approaches face inherent performance limitations when constrained to this single-round interaction paradigm, as illustrated in Figure 2. The top part visualizes the RL methods modeling the interactions with single user as Markov processes. In contrast, the bottom part shows Two-stage and DFL methods making decisions based on current state alone. In these settings, the potential for leveraging the richer sequential dependencies in user behavior is significantly curtailed, leading to a performance plateau.

Unlike these single-round paradigms, our work, while inheriting the benefits of an end-to-end framework, embraces the sequential nature of user engagement within e-commerce platforms. We propose an efficient framework named **Sequence-Aware Constrained Optimization (SACO)** framework to directly devise coupon distribution policy for long-term revenue boosting. Critically, while demonstrated here for coupon distribution, SACO’s general framework empowers exploration of constrained optimization scenarios across many applications. Our contributions are threefold:

Generalized and Practical Problem Setting. We address more complex scenarios where: (1) multiple users arrive simultaneously per round (e.g., group-buying scenarios), requiring parallel coupon allocation; and (2) users may revisit the platform across multiple rounds, creating sequential dependencies.

Efficient Framework. SACO framework offers two key advantages: (1) by explicitly modeling the series of user in-

teractions and utilizing historical e-commerce data, SACO captures valuable temporal dynamics overlooked by single-round approaches; and (2) by reformulating sequential decision-making as an auto-regressive prediction task, we leverage causal-Transformer structure to effectively capturing multi-round causal relationships. Meanwhile, this structure reduces the framework’s reliance on the statistical premise of unbiased data, which serves as the training set foundation for two-stage methods and DFL based on uplift models. This allows for efficient, timely, and cost-effective iterative updates, contributing to sustained optimal performance in dynamic online environments.

Extensive Experiments on Public, Synthetic and Industrial Datasets. We validate our approach through comprehensive experiments on diverse datasets. Comparisons against state-of-the-art baselines demonstrate an average revenue improvement of **3.60%**. Supporting experiments include ablation studies and inference speed analysis.

Further Related Work

Coupon distribution requires estimating user responses to different coupon strategies, a task known as counterfactual evaluation. A key metric is the conditional average treatment effect (CATE), which quantifies the impact of a specific coupon on outcomes like conversion rates, considering individual user characteristics (Devriendt, Moldovan, and Verbeke 2018). Uplift modeling, a field dedicated to CATE estimation, has produced various machine learning techniques, often trained on randomized controlled trial (RCT) data (Athey and Imbens 2015). The trained models then predict individual-level causal effects for arbitrary interventions. Notable CATE estimation techniques include meta-learners (Hansotia and Rukstales 2002; Athey and Imbens 2015; Künzel et al. 2019; Nie and Wager 2021), ranking-based methods (Kuusisto et al. 2014; Betlei, Diemert, and Amini 2021), representation learning (Johansson, Shalit, and Sontag 2016; Shi, Blei, and Veitch 2019; Yao et al. 2018), and causal forests (Ai et al. 2022; ATHEY, TIBSHIRANI, and WAGER 2019; Wager and Athey 2018; Zhao, Fang, and Simchi-Levi 2017). However, existing CATE formulations primarily address single-treatment effects. In contrast, our problem requires modeling *temporal treatment dependencies*, where user responses depend not only on the current coupon but also on historical coupon exposures. This introduces complex sequential counterfactual estimation challenges that traditional uplift models cannot address.

To tackle this challenge, we leverage the Transformer architecture (Vaswani 2017), which has demonstrated remarkable capabilities in sequential modeling across diverse domains including natural language processing (Keskar et al. 2019; Dai et al. 2019), computer vision (Parvaiz et al. 2023), and reinforcement learning (RL) (Chen et al. 2021; Zheng, Zhang, and Grover 2022; Wu, Wang, and Hamaya 2023). While prior RL applications focus on unconstrained optimization of cumulative rewards, we adapt the Decision Transformer (DT) framework (Chen et al. 2021) for constrained optimization problems. Our modification enables handling of budget constraints while preserving DT’s ability to model sequential dependencies in decision-making.

Coupon Distribution Problem

Model on Coupon Distribution

We formulate the coupon distribution problem as a budget allocation problem with multiple treatments, which was also introduced in many existing works. Different from them, we consider a comprehensive scenario with sequential coupon distribution decisions for diverse users who continuously interact with the e-commerce platform.

An e-commerce platform has N users. Given the requirement that all users register on the platform before purchasing, N is a known parameter. Over a finite time horizon, users interact with the platform for up to T times, referred to as T rounds. Upon receiving a user interaction request, the platform must make a real-time decision regarding coupon issuance. Critically, once a coupon is granted, it cannot be revoked. Over the T rounds, the platform has a marketing budget B , which is the total amount it is willing to spend on issuing coupons to all users. The platform offers K types of coupons, such as $\{2, 5, 10\}$ dollar coupons for $K = 3$, and so forth. Let $[N], [T], [K]$ denote the sets of users, rounds, and coupons, respectively.

We define x_{itk} as the platform’s coupon issuance decision, where $x_{itk} = 1$ indicates that in round t , the platform issues coupon k to user i , with $i \in [N], t \in [T], k \in [K]$. Following coupon issuance, a user generates consumption r_{itk} , while the platform incurs a corresponding cost of c_{itk} as a subsidy. The platform’s objective is to maximize revenue from marketing within the limited budget B . Formally, we perform a linear relaxation on x_{itk} and express the platform’s online marketing problem as the following online optimization problem:

$$\max \sum_{i \in [N], t \in [T], k \in [K]} r_{itk} x_{itk} \quad (1)$$

$$\text{s.t.} \quad \sum_{i \in [N], t \in [T], k \in [K]} c_{itk} x_{itk} \leq B, \quad (2)$$

$$0 \leq x_{itk} \leq 1, \forall i \in [N], t \in [T], k \in [K]. \quad (3)$$

In previous research, the platform receives a request from only one user per round. In this paper, we generalize this constraint by allowing the platform to receive requests from multiple users per round. For instance, in group-buying scenarios, the platform needs to make coupon issuance decisions for multiple users simultaneously. Furthermore, by leveraging Transformer architectures to process multi-dimensional action space, we allow the platform to issue a bundle of coupons (e.g. the combination of 5-dollar coupon and 10% discount) to each user per time, which aligns more closely with real-world complex scenarios.

Dual Problem Formulation

Let \mathbf{x} represent the coupon distribution strategy, $R(\mathbf{x})$ denote the total revenue obtained by the platform under this strategy, and $C(\mathbf{x})$ represent the total amount of coupons subsidized by the platform under this strategy. To obtain the optimal distribution strategy within the budget constraint, we first introduce λ as the dual variable for the constraint

$\sum_{i \in [N], t \in [T], k \in [K]} c_{itk} x_{itk} \leq B$. Next, we define the Lagrangian function $L(\mathbf{x}, \lambda)$ and the objective function of the dual problem is the maximum value of the Lagrangian function with respect to \mathbf{x} , i.e.,

$$\min_{\lambda \geq 0} \max_{0 \leq x_{itk} \leq 1} L(\mathbf{x}, \lambda) \quad (4)$$

$$= \min_{\lambda \geq 0} \max_{0 \leq x_{itk} \leq 1} (R(\mathbf{x}) - \lambda (C(\mathbf{x}) - B)) \quad (5)$$

Problem (1) is formulated as a linear program. Therefore, strong duality holds, implying that we can find the optimal primal solution by solving the dual problem. Given optimal dual variable λ^* and known parameter B , the optimization objective simplifies to

$$\max_{0 \leq x_{itk} \leq 1} (R(\mathbf{x}) - \lambda^* C(\mathbf{x})). \quad (6)$$

Assuming all responses are known, the problem simplifies to a classical multiple choice knapsack problem, yielding an approximation ratio near 1 (Zhou et al. 2023, 2024). However, the platform can only observe the actual revenue or cost after a specific coupon decision and the counterfactual data is inherently unattainable. Furthermore, user responses are sequential. Previous coupon distribution decisions influence subsequent users’ coupon responses, making counterfactual prediction challenging. Therefore, the online linear programming problem cannot be directly solved.

Our Framework

To address the above challenge, we propose our framework named **Sequence-Aware Constrained Optimization (SACO)** framework, a direct policy learning approach involving training a model to determine the optimal sequential policy \mathbf{x}_i for each user i given λ , then leveraging a λ optimization algorithm and model-parallel inference to derive the optimal overall strategy \mathbf{x} .

Trajectory Representation

In the e-commerce platform, learning policies through direct interaction with the online environment is infeasible. However, by organizing historical online logs into structured trajectories, we unlock a feasible strategy for effective model training. We model the sequential decision-making problem for single-user and propose the trajectory representation as the training dataset format.

The sequential decision-making problem is represented by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \mathcal{C}, \gamma \rangle$. Specifically,

- \mathcal{S} : The state space, which describes the user context, including user features and other relevant information. $s_t \in \mathcal{S}$ represents the user context for all users in round t , where $\mathbf{s}_t = (s_{1t}, \dots, s_{it}, \dots, s_{Nt})$.
- \mathcal{A} : The coupon distribution action space, which encompasses the choices of treatments for users. $\mathbf{a}_t \in \mathcal{A}$ represents the coupon distribution strategy in round t , where $\mathbf{a}_t = (x_{1t1}, \dots, x_{Nt1}, \dots, x_{itk}, \dots, x_{NtK})$.
- \mathcal{P} : The state transition function, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$. The transition dynamics $P(s_{t+1} | s_t, a_t)$ describe the dynamic transition from the current state s_t to the next state s_{t+1} after taking an action a_t .

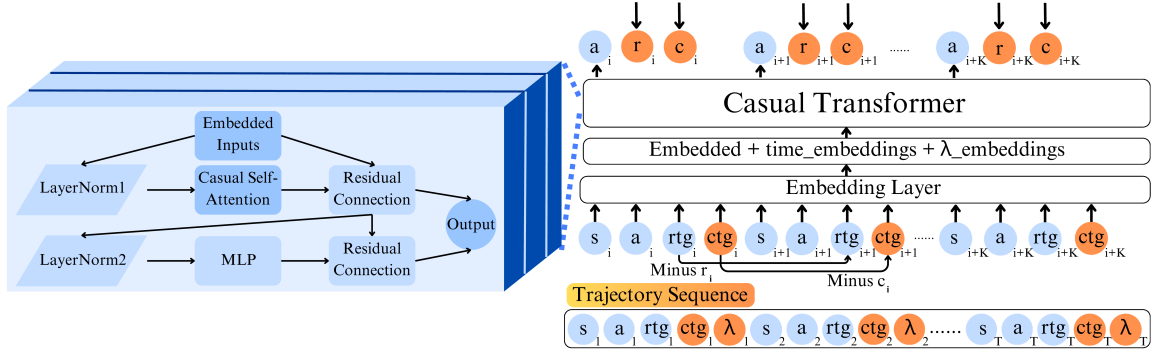


Figure 3: Constrained-Optimization Decision Transformer

- \mathcal{R} : The reward space, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R} \in R$, which describes the platform revenue generated by coupon distribution. $r_t \in \mathcal{R}$ represents the platform revenue in round t , where $r_t = \sum_{i \in [N]} r_{it}$ and $r_{it} = \sum_{k \in [K]} r_{itk}$.
- \mathcal{C} : The cost space, $\mathcal{C} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{C} \in R$, which describes the marketing budget expenditure incurred by coupon distribution actions. $c_t \in \mathcal{C}$ represents the marketing budget expenditure in round t , where $c_t = \sum_{i \in [N]} c_{it}$ and $c_{it} = \sum_{k \in [K]} c_{itk}$.
- γ : The discount factor for future rewards, indicating the importance of future rewards relative to the present. In this paper, we set $\gamma = 1$.

Given a dataset in the format (user-id, user-features as state, treatment as action, cost, reward, time) and additional λ randomized in the interval $(0, 1)$, we aggregate the dataset by user-id and then sort the data with each user's group in chronological order. We define $RTG_{it} = \sum_{t'=t}^T \gamma r_{it}$ (Chen et al. 2021) and $CTG_{it} = \sum_{t'=t}^T c_{it}$. The training dataset is pre-processed into trajectory data of size N :

$$\tau_{i \in [N]} = [(s_{i_0}, a_{i_0}, CTG_{i_0}, RTG_{i_0}, \lambda_i), \dots, (s_{i_t}, a_{i_t}, CTG_{i_t}, RTG_{i_t}, \lambda_i), \dots, (s_{i_T}, a_{i_T}, CTG_{i_T}, RTG_{i_T}, \lambda_i)]$$

As mentioned above, the dual variable λ reflects the constraint on the budget B . To improve the model's ability to generalize across diverse budget levels, we randomize λ ten times for each trajectory sample. This expands our training dataset tenfold, promoting more robust learning.

Constrained-Optimization Decision Transformer

Having established the training dataset, we adapt the Decision Transformer model and align it for direct policy learning, as summarized in Figure 3 and Algorithm 1.

Our model leverages two key advantages of the causal transformer structure, the ability to handle sequential data effectively via its transformer architecture and the capability to evaluate counterfactual outcomes through its causal structure (Nichani, Damian, and Lee 2024).

To integrate constraint information, we augment the input space with the parameters ctg , representing constraint

Algorithm 1: Model Training

Input: State dimension d_s , action dimension d_a
Maximum episode length T , sequence length K
Parameter: Learning rates η , weight decay β

Output: Trained model θ

- 1: **Initialize:** θ , AdamW(θ, η, β)
- 2: **for** epoch = 1 **to** N_{epochs} **do**
- 3: Sample batch $\mathcal{B} = \{(s_t, a_t, rtg_t, ctg_t, \lambda)\}_{t=j}^{j+K}$
- 4: Normalize states: $s_t \leftarrow (s_t - \mu_s) / \sigma_s$
- 5: **Compute embeddings:**
- 6: $e_s \leftarrow \text{Embed}(s_t) + \text{Embed}(t) + \text{Embed}(\lambda_t)$
- 7: $e_a \leftarrow \text{Embed}(a_t) + \text{Embed}(t) + \text{Embed}(\lambda_t)$
- 8: $e_{RTG} \leftarrow \text{Embed}(RTG_t) + \text{Embed}(t) + \text{Embed}(\lambda_t)$
- 9: $e_{CTG} \leftarrow \text{Embed}(CTG_t) + \text{Embed}(t) + \text{Embed}(\lambda_t)$
- 10: Stack embeddings: $x \leftarrow \text{Stack}(e_s, e_a, e_{RTG}, e_{CTG})$
- 11: Apply layer norm: $x \leftarrow \text{LayerNorm}(x)$
- 12: **for each transformer block** Block_i **do**
- 13: $x \leftarrow \text{Block}_i(x)$
- 14: **end for**
- 15: **Predict:** \hat{a}_t
- 16: **Compute task losses:** \mathcal{L}_a
- 17: Backpropagate \mathcal{L}_a and update θ using optimizers
- 18: **end for**
- 19: **return** Trained model θ

types. Furthermore, we enhance the original time embedding $\text{Embed}(t)$ by incorporating the dual variable embedding $\text{Embed}(\lambda)$, as shown in line 6-9 of Algorithm 1. This enhancement effectively encoded constraint information alongside temporal signals for each parameter. Subsequently, these augmented inputs are processed through the Causal Transformer blocks, with the internal details visualized in the left panel of Figure 3.

As shown in line 16 of Algorithm 1, we employed a cross-entropy loss calculated on the action predictions:

$$\mathcal{L}_a(\theta) = -\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \sum_{j=1}^K y_{itj} \log(p_\theta(a_{it} = j | s_{it})) \quad (7)$$

only when $j = j'$, then $y_{itj} = 1$, meaning that the j' type of coupon is distributed to the user in the dataset and the

Algorithm 2: Overall Algorithm of SACO

Input: N , user states s , budget B
Parameter: max_iterations, noise, learning_rate lr
Output: λ^* , R^*

- 1: **Initialize:** $\lambda \leftarrow \text{Uniform}(0, 1)$
- 2: **Define Function OBJECTIVE(λ):**
- 3: Simulate all users with λ through trained model
- 4: Compute total cost C and total revenue R
- 5: penalty $\leftarrow \lambda \cdot \max(C(x) - B, 0)$
- 6: **return** penalty $- R(x)$
- 7: **while** $C > B$ or $C < B - \epsilon$ **do**
- 8: Reset the offline simulation environment
- 9: init $\leftarrow \lambda + \text{Uniform}(-\text{noise}, \text{noise})$
- 10: init $\leftarrow \text{Clip}(\text{init}, 0, 1)$
- 11: **Optimize** λ **using SLSQP:**
- 12: result $\leftarrow \text{minimize}(\text{OBJECTIVE}, \text{init}, \text{SLSQP}, \text{bounds})$
- 13: $\lambda \leftarrow \text{result}.x[0]$ {Update λ }
- 14: $R, C \leftarrow (\text{user_states}, \lambda)$
- 15: **if** $R > R^*$ **and** $C \leq B$ **then**
- 16: $R^* \leftarrow R, \lambda^* \leftarrow \lambda$
- 17: **end if**
- 18: $\lambda \leftarrow \lambda + lr \cdot (C - B)$ {Dynamic adjustment}
- 19: **end while**
- 20: **return** λ^*, R^*

response is recorded and observable. This design choice enables the model to learn optimal policies tailored to specific optimization objectives and constraints.

Following training, we obtained the model capable of making optimal sequential decisions for individual users, conditioned on the dual variables.

Overall Strategy

Based on the trained model to determine the optimal sequential policy x_i for each user i given λ , we combine it with the λ optimization algorithm, obtaining the SACO framework to derive the optimal overall strategy x , as summarized in Algorithm 2.

We adopt the SLSQP algorithm to optimize Lagrangian multiplier λ , as it is versatile and robust. SLSQP eliminates the need for fine-tuning while delivering reliable solutions for practical engineering applications.

We set the objective of SLSQP algorithm as

$$\lambda \cdot (C(x) - B) - R(x). \quad (8)$$

Considering the limitation that the SLSQP algorithm may fall into the local optimum, we will perform multiple rounds of iterative optimization in our framework. In each new round of iterative optimization process, we will add some noise disturbance and penalty on λ to appropriately guide the update direction, shown in line 9 and 18 of Algorithm 2.

Assumption 1 When $C(x) > C(x')$, then $R(x) > R(x')$.

We operate under the assumption of a direct correlation between reward magnitude and associated cost, an observation validated by typical marketing and advertising practices. Hence, we define the convergence criterion for iteration as $B - \epsilon < C < B$.

Finally, we obtain the overall framework to directly devise coupon distribution policy for long-term revenue boosting.

Unlike prior single-round approaches, SACO, as a sequential modeling framework, unlocks the potential to extract complex user-platform interaction patterns from rich datasets. Critically, while demonstrated here for coupon distribution, SACO’s general framework empowers exploration of constrained optimization scenarios across a spectrum of applications. Furthermore, its significant scalability allows for straightforward integration of external components.

Experiment

In this section, we conduct experiments to evaluate the performance of our proposed framework. Specifically, we compare it with baseline algorithms, perform an ablation study to assess the contribution of each component, and conduct an inference speed analysis. Since both our framework and the baselines may generate actions \tilde{a} that are not present in the dataset trajectories, the corresponding revenue \tilde{r} and cost \tilde{c} for these actions are missing. To address this issue of missing data, we follow the approach outlined in (Guo et al. 2024) and create an offline simulation environment based on the tuples (s, a, r, c) within the existing trajectories. This environment enables us to obtain the tuples $(s, \tilde{a}, \tilde{r}, \tilde{c})$ for actions that are not included in the original trajectories.

Experimental Settings

Dataset To demonstrate the robustness of our framework, we leverage a diverse set of datasets including public benchmarks, real-world industrial data, and synthetic simulations.

CRITEO-UPLIFT v2 (CU). The dataset is provided by Criteo AI Labs at the AdKDD’18 workshop (Diemert et al. 2018) and comprises 13.9 million samples collected from Randomized Controlled Trial (RCT). Each sample contains 12 features, a binary treatment indicator, and two response labels, namely visit and conversion. We follow the approach in (Zhou et al. 2023), extracting user feature columns ‘f0-f11’ as states, ‘treatment’ as actions, ‘visit’ label as costs and ‘conversion’ label as generated revenue.

MT-LIFT (MT). Collected from two months of coupon marketing scenarios for food delivery(Huang et al. 2024), the dataset contains nearly 5.5 million samples, which include 99 features, 5 types of treatment, and two response labels: click and conversion. Similar to the CU dataset, we process user feature columns ‘f0-f98’ as states, ‘treatment’ as actions, ‘click’ as costs and ‘conversion’ as revenue.

Synthetic Dataset (SY). We have developed a synthetic dataset generation pipeline inspired by (Yan et al. 2023). The dataset consists of 10,000 simulated users. Each user interaction spans 10 timesteps, with an action space of four treatments applied randomly with equal probabilities at each timestep. The cost and reward is generated when the user responds to the coupon with the response probability. Gaussian noise is used to simulate state transitions.

Business Dataset (BD). To realistically evaluate the performance of our model in real-world industrial scenarios, we incorporate a recent business dataset originating from ByteDance, a global leading technology company with

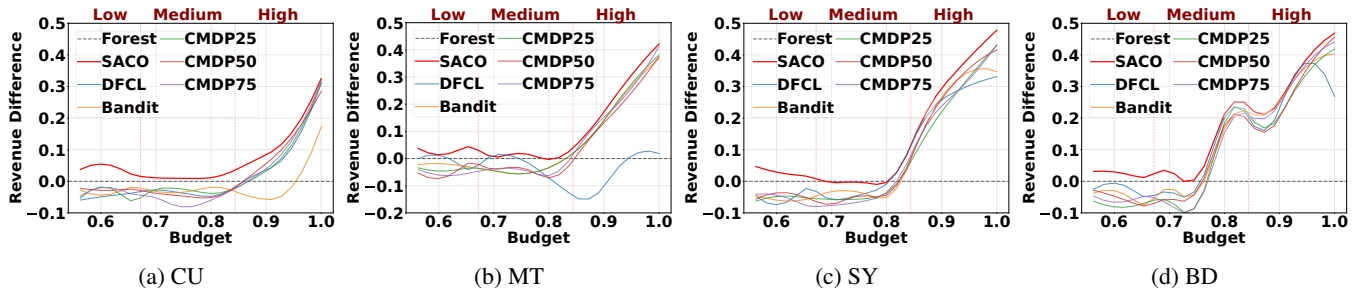


Figure 4: Evaluation Results

Dataset	Metric	SACO	DFCL	Bandit	CMDP25	CMDP50	CMDP75	Forest
CU	\overline{ROI}	1.11	1.052	1.072	1.053	1.053	1.052	1.087
	\overline{BARate}	0.997	0.986	0.966	0.987	0.991	0.984	0.968
MT	\overline{ROI}	1.105	1.106	1.06	1.053	1.054	1.059	1.088
	\overline{BARate}	1.0	0.926	0.988	0.989	0.984	0.984	0.952
SY	\overline{ROI}	1.109	1.073	1.066	1.053	1.053	1.057	1.101
	\overline{BARate}	0.999	0.977	0.978	0.983	0.986	0.98	0.934
BD	\overline{ROI}	1.126	1.069	1.067	1.05	1.057	1.052	1.101
	\overline{BARate}	0.996	0.981	0.983	0.979	0.984	0.985	0.919

Table 1: Performance Comparison Across Datasets and Baselines

a range of popular platforms. User features included are anonymized to ensure data privacy and compliance. The marketing dataset consists of millions of samples, including anonymized user features as states, treatments as actions, ‘cost’ label as costs and ‘is paid’ label as revenue.

For all datasets, we randomly split them separately, using 80% of the samples for training and 20% for testing. After initial data cleaning to remove anomalies, we aggregate the data by user ID or relatively stable user feature. Within each aggregated group, we subsequently sort the data chronologically to generate trajectory samples.

Evaluation Metrics

- Revenue ($R(\mathbf{x})$): A crucial factor in measuring marketing effectiveness is the platform revenue generated by the marketing efforts.
- Return on Investment (ROI) ($\frac{R(\mathbf{x})}{C(\mathbf{x})}$): The emphasis on the ROI indicator is paramount. Without proper consideration of cost, the strategy risks favoring marketing activities targeting users with both high returns and high associated expenditures. It leads to a suboptimal outcome and inefficient resource allocation.
- Budget Allocation Rate (BARate) ($\frac{C(\mathbf{x})}{B}$): Effective coupon distribution requires assessing user sensitivity beforehand. Successful strategies will exhibit high budget utilization, implying that coupons are effectively targeted at highly sensitive users.

Evaluation Experiments

Baseline Algorithms

- Causal **Forest**(Athey and Wager 2021): This baseline model, commonly used in marketing scenarios, is trained

on historical data to predict individual treatment effects (ITE). At each decision point, the model estimates the treatment effect of offering a coupon and allocates it only if the predicted effect is positive.

- **Bandit** (Chapelle and Li 2011): We employ a multi-arm bandit algorithm based on Thompson Sampling (TS). This algorithm leverages Bayesian methods to construct a belief over the reward distribution of actions.
- **Constrained Markov Decision Process (CMDP)** (Xiao et al. 2019): For different λ values, the CMDP model needs to be retrained. Hence, we established three CMDP baselines, **CMDP25**, **CMDP50** and **CMDP75**, by re-training the core policy network with λ values of 0.25, 0.5, and 0.75, respectively.
- **Decision Focused Causal Learning (DFCL)** (Zhou et al. 2024): The DFCL baseline is a state-of-the-art method that claims to achieve breakthrough performance by building upon two-stage approaches. It is an end-to-end framework, aiming to align the prediction objective with the ultimate decision-making goal.

Experimental Results To establish a clear baseline, we select the Forest algorithm as our benchmark model. Figure 4 presents the revenue difference between evaluated models and the causal forest benchmark. We choose Forest due to its widespread adoption in marketing platforms, known for its robust causal learning capabilities. It is an appropriate and meaningful point of comparison.

To assess budget constraint robustness, we defined three budget levels: Low, Medium, and High, representing strict to loose constraints. Within each level, we specify three distinct budget amounts. All budget values are normalized to avoid dependence on specific numerical scales and ensure




Figure 5: Ablation Study on MT Dataset.

generalizability across different scenarios.

The experimental results in Figure 4 demonstrate that our proposed SACO framework gains higher revenue compared to baselines across all datasets and budget levels, with two exceptions: the medium budget level in the SY dataset and MT dataset. SACO achieves an average revenue improvement of **3.60%**. Furthermore, Table 1 highlights SACO’s superior performance in terms of both average ROI and BARate across the tested budget range. This overall superior performance emphasizes SACO’s strong strategy learning capabilities, robust causal reasoning, and adaptability to diverse dataset characteristics.

We hypothesize that SACO’s slightly lower revenue than DFCL at the medium budget level in MT is because DFCL, developed by Meituan, may be particularly suited to the Meituan dataset. Regarding the slightly lower revenue compared to Forest at the medium budget level in the SY dataset, this is attributed to the uniform distribution of actions within that specific synthetic dataset. It is worth noting that forest and uplift models often benefit from training on unbiased RCT data, which features a uniform distribution of actions. However, online deployment would typically utilize an optimized strategy like DFCL algorithm, rather than a random coupon distribution, which means the newly sampled data inherently introduces selection bias. While some existing methods don’t detail how their models would be iteratively updated under these biased conditions, our SACO’s implicit causal learning approach offers a distinct advantage. By reducing its reliance on the statistical premise of unbiased data, SACO allows for efficient, timely, and cost-effective iterative updates, contributing to sustained optimal performance in dynamic online environments.

Ablation Study

Ablation Design

- **No-constraint:** We omit the parameter λ and \widehat{RTC} from our model, assessing their significance on learning and respecting budget constraints.
- **No- \widehat{RTG} :** We utilize the original Transformer model (Vaswani 2017) without the \widehat{RTG} parameter to evaluate its impact on the model’s capacity to pursue long-term revenue, as opposed to immediate rewards.

Experimental Results We conduct ablation studies in the MT-LIFT public dataset. The experimental results are presented in Figure 5. The horizontal axis depicts revenue (positive values) and costs (negative values), while the vertical

axis lists the ablated algorithms. The dashed vertical line denotes the budget constraint.

Comparing our framework with No-constraint model reveals that removing constraint-related variables leads to a diminished awareness of budgetary limitations. The model fails to adequately respond to budgetary changes. Comparing SACO to model without RTG shows that, when the budget is not sufficiently generous, the absence of RTG causes the model to marginally exceed the budget in pursuit of immediate revenue gains. Furthermore, when budget is ample, the absence of RTG guidance restricts overall revenue generation, rendering its performance inferior to SACO.

Inference Speed Analysis

Motivated by the demands of real-world applications, we conduct inference efficiency experiments to validate the runtime performance of SACO. The results are presented below.

Scenarios	Time (ms)	Hardware
Single-user	24.6243	CPU-only
Single-user	3.8128	NVIDIA RTX 4090
128 users	48.5075	NVIDIA RTX 4090

Table 2: Model Inference Time

Based on the results, SACO framework achieves an inference speed suitable for industrial applications. In addition, we use 16 parallel models in the multi-user scenario, limited by CUDA memory. We anticipate that deployment in industrial environments, equipped with superior computing resources (e.g., larger GPU memory, parallel GPUs), will effectively mitigate this memory bottleneck.

Conclusion

In this paper, we propose SACO framework, sequentially formulating coupon distribution strategies to enhance long-term revenue. Our unified framework combines general scenarios, sequential modeling and efficient iterative updates. Extensive experiments have been conducted to demonstrate the superiority of our framework.

Future directions includes integrating SACO with externalities, such as incorporating user-to-user impact evaluation as input features. Also, we plan to encode coupons to prevent action space drift. This is in case we need to dynamically adapt coupon types (e.g., larger discounts) based on observed coupon performance.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No.62472428), Public Computing Cloud, Renmin University of China, the fund for building world-class universities (disciplines) of Renmin University of China and ByteDance.

References

- Ai, M.; Li, B.; Gong, H.; Yu, Q.; Xue, S.; Zhang, Y.; Zhang, Y.; and Jiang, P. 2022. Lbcf: A large-scale budget-constrained causal forest algorithm. In *Proceedings of the ACM Web Conference 2022*, 2310–2319.
- Albert, J.; and Goldenberg, D. 2022. E-commerce promotions personalization via online multiple-choice knapsack with uplift modeling. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2863–2872.
- Athey, S.; and Imbens, G. W. 2015. Machine learning methods for estimating heterogeneous causal effects. *stat*, 1050(5): 1–26.
- ATHEY, S.; TIBSHIRANI, J.; and WAGER, S. 2019. GENERALIZED RANDOM FORESTS. *The Annals of Statistics*, 47(2): 1148–1178.
- Athey, S.; and Wager, S. 2021. Policy learning with observational data. *Econometrica*, 89(1): 133–161.
- Betlei, A.; Diemert, E.; and Amini, M.-R. 2021. Uplift modeling with generalization guarantees. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 55–65.
- Chapelle, O.; and Li, L. 2011. An empirical evaluation of thompson sampling. *Advances in neural information processing systems*, 24.
- Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q. V.; and Salakhutdinov, R. 2019. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.
- Devriendt, F.; Moldovan, D.; and Verbeke, W. 2018. A literature survey and experimental evaluation of the state-of-the-art in uplift modeling: A stepping stone toward the development of prescriptive analytics. *Big data*, 6(1): 13–41.
- Diemert, E.; Betlei, A.; Renaudin, C.; and Amini, M.-R. 2018. A large scale benchmark for uplift modeling. In *KDD*.
- Du, S.; Lee, J.; and Ghaffarizadeh, F. 2019. Improve user retention with causal learning. In *The 2019 ACM SIGKDD Workshop on Causal Discovery*, 34–49. PMLR.
- Goldenberg, D.; Albert, J.; Bernardi, L.; and Estevez, P. 2020. Free lunch! retrospective uplift modeling for dynamic promotions recommendation within roi constraints. In *Proceedings of the 14th ACM Conference on Recommender Systems*, 486–491.
- Guo, J.; Huo, Y.; Zhang, Z.; Wang, T.; Yu, C.; Xu, J.; Zheng, B.; and Zhang, Y. 2024. Generative Auto-bidding via Conditional Diffusion Modeling. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5038–5049.
- Hansotia, B. J.; and Rukstales, B. 2002. Direct marketing for multichannel retailers: Issues, challenges and solutions. *Journal of Database Marketing & Customer Strategy Management*, 9: 259–266.
- Huang, Y.; Wang, S.; Gao, M.; Wei, X.; Li, C.; Luo, C.; Zhu, Y.; Xiao, X.; and Luo, Y. 2024. Entire Chain Uplift Modeling with Context-Enhanced Learning for Intelligent Marketing. In *Companion Proceedings of the ACM on Web Conference 2024*, 226–234.
- Joachims, T.; Swaminathan, A.; and De Rijke, M. 2018. Deep learning with logged bandit feedback. In *International Conference on Learning Representations*.
- Johansson, F.; Shalit, U.; and Sontag, D. 2016. Learning representations for counterfactual inference. In *International conference on machine learning*, 3020–3029. PMLR.
- Keskar, N. S.; McCann, B.; Varshney, L. R.; Xiong, C.; and Socher, R. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10): 4156–4165.
- Kuusisto, F.; Costa, V. S.; Nassif, H.; Burnside, E.; Page, D.; and Shavlik, J. 2014. Support vector machines for differential prediction. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part II 14*, 50–65. Springer.
- Nichani, E.; Damian, A.; and Lee, J. D. 2024. How transformers learn causal structure with gradient descent. *arXiv preprint arXiv:2402.14735*.
- Nie, X.; and Wager, S. 2021. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2): 299–319.
- Parvaiz, A.; Khalid, M. A.; Zafar, R.; Ameer, H.; Ali, M.; and Fraz, M. M. 2023. Vision Transformers in medical computer vision—A contemplative retrospection. *Engineering Applications of Artificial Intelligence*, 122: 106126.
- Shi, C.; Blei, D.; and Veitch, V. 2019. Adapting neural networks for the estimation of treatment effects. *Advances in neural information processing systems*, 32.
- Swaminathan, A.; and Joachims, T. 2015a. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1): 1731–1755.
- Swaminathan, A.; and Joachims, T. 2015b. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, 814–823. PMLR.

Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.

Wager, S.; and Athey, S. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523): 1228–1242.

Wu, Y.-H.; Wang, X.; and Hamaya, M. 2023. Elastic decision transformer. *Advances in neural information processing systems*, 36: 18532–18550.

Xiao, S.; Guo, L.; Jiang, Z.; Lv, L.; Chen, Y.; Zhu, J.; and Yang, S. 2019. Model-based constrained MDP for budget allocation in sequential incentive marketing. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 971–980.

Yan, Z.; Wang, S.; Zhou, G.; Lin, J.; and Jiang, P. 2023. An end-to-end framework for marketing effectiveness optimization under budget constraint. *arXiv preprint arXiv:2302.04477*.

Yao, L.; Li, S.; Li, Y.; Huai, M.; Gao, J.; and Zhang, A. 2018. Representation learning for treatment effect estimation from observational data. *Advances in neural information processing systems*, 31.

Ye, P.; Qian, J.; Chen, J.; Wu, C.-h.; Zhou, Y.; De Mars, S.; Yang, F.; and Zhang, L. 2018. Customized regression model for airbnb dynamic pricing. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 932–940.

Zhang, Y.; Tang, B.; Yang, Q.; An, D.; Tang, H.; Xi, C.; Li, X.; and Xiong, F. 2021. BCORLE (λ): An Offline Reinforcement Learning and Evaluation Framework for Coupons Allocation in E-commerce Market. *Advances in Neural Information Processing Systems*, 34: 20410–20422.

Zhao, K.; Hua, J.; Yan, L.; Zhang, Q.; Xu, H.; and Yang, C. 2019. A unified framework for marketing budget allocation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1820–1830.

Zhao, Y.; Fang, X.; and Simchi-Levi, D. 2017. Uplift modeling with multiple treatments and general response types. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, 588–596. SIAM.

Zheng, Q.; Zhang, A.; and Grover, A. 2022. Online decision transformer. In *international conference on machine learning*, 27042–27059. PMLR.

Zhou, H.; Huang, R.; Li, S.; Jiang, G.; Zheng, J.; Cheng, B.; and Lin, W. 2024. Decision focused causal learning for direct counterfactual marketing optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 6368–6379.

Zhou, H.; Li, S.; Jiang, G.; Zheng, J.; and Wang, D. 2023. Direct heterogeneous causal learning for resource allocation problems in marketing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 5446–5454.