

# Direct Heterogeneous Causal Learning for Resource Allocation Problems in Marketing

Hao Zhou<sup>1\*</sup>, Shaoming Li<sup>1</sup>, Guibin Jiang<sup>1</sup>, Jiaqi Zheng<sup>2</sup>, Dong Wang<sup>1</sup>

<sup>1</sup> Meituan, China

<sup>2</sup> State Key Laboratory for Novel Software Technology, Nanjing University, China  
{zhouhao29, lishaoming, jiangguibin, wangdong07}@meituan.com, jzheng@nju.edu.cn

## Abstract

Marketing is an important mechanism to increase user engagement and improve platform revenue, and heterogeneous causal learning can help develop more effective strategies. Most decision-making problems in marketing can be formulated as resource allocation problems and have been studied for decades. Existing works usually divide the solution procedure into two fully decoupled stages, i.e., machine learning (ML) and operation research (OR) — the first stage predicts the model parameters and they are fed to the optimization in the second stage. However, the error of the predicted parameters in ML cannot be respected and a series of complex mathematical operations in OR lead to the increased accumulative errors. Essentially, the improved precision on the prediction parameters may not have a positive correlation on the final solution due to the side-effect from the decoupled design.

In this paper, we propose a novel approach for solving resource allocation problems to mitigate the side-effects. Our key intuition is that we introduce the decision factor to establish a bridge between ML and OR such that the solution can be directly obtained in OR by only performing the sorting or comparison operations on the decision factor. Furthermore, we design a customized loss function that can conduct direct heterogeneous causal learning on the decision factor, an unbiased estimation of which can be guaranteed when the loss converges. As a case study, we apply our approach to two crucial problems in marketing: the binary treatment assignment problem and the budget allocation problem with multiple treatments. Both large-scale simulations and online A/B Tests demonstrate that our approach achieves significant improvement compared with state-of-the-art.

## Introduction

Marketing is one of the most effective mechanisms for the improvement of user engagement and platform revenue. Thus, various marketing campaigns have been widely deployed in many online Internet platforms. For example, markdowns of perishable products in Freshippo are used to boost sales (Hua et al. 2021), coupons in Taobao Deals can stimulate user activity (Zhang et al. 2021) and incentives in the Kuaishou video platform are offered to improve user retention (Ai et al. 2022).

\*Corresponding author

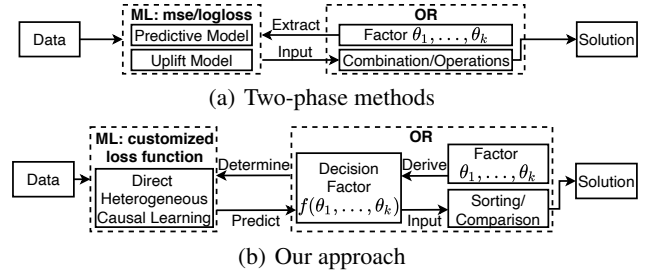


Figure 1: Comparison with two-phase methods

Despite the incremental revenue, marketing actions can also incur significant consumption of marketing resources (e.g., budget). Hence, only part of individuals (e.g., shops or goods) may be assigned the marketing treatments due to the limited volume of them. In marketing, such decision-making problems can be formulated as resource allocation problems, and have been investigated for decades.

Most of existing studies usually use two-phase methods to solve these problems (Ai et al. 2022; Zhao et al. 2019; Du, Lee, and Ghaffarizadeh 2019). As is shown by Fig. 1(a), the first stage is ML, where the (incremental) response of individuals under different treatments is predicted by predictive/uplift models. The second stage is OR and the prediction results in ML are fed as the input to the combinatorial optimization algorithms. Hence, existing works mainly focus on the decoupled optimization for predictive/uplift modeling and combinatorial optimization.

Despite the widespread application, there are two major defects in two-phase methods. The first one is that the solution is obtained after conducting many intermediate computations on the prediction results in ML, e.g., the combination of multiple factors or complex mathematical operations in OR. Therefore, the improved precision on the prediction parameters may not have a positive correlation on the final solution. The second is that the errors of model prediction are not respected, and the complex operations performed on prediction results in OR lead to the increased accumulative errors. Due to the accumulative errors, the theoretically optimal algorithm in OR cannot always achieve the practically optimum and is even inferior to heuristic strategies in some scenarios. Therefore, the decoupled optimization for

ML and OR cannot induce a global optimization for the original problem.

Instead of two-phase methods, we propose a novel approach for solving resource allocation problems to mitigate the above defects. First of all, we define the decision factor of an algorithm as a factor on which the solution can be directly obtained by performing only the sorting or comparison operations. As is shown by Fig. 1(b), the decision factor derived from OR is taken as the learning objective, and direct heterogeneous causal learning in ML is conducted in our approach. Based on the definition, there is no alternative mathematical operations on the prediction results in OR. Therefore, the ranking performance of one model on the decision factor directly determines the quality of the solution and improving the model can guarantee a better solution. Specifically, the model error can be used to measure the ranking performance instead, which is respected and not enlarged in OR. Hence, the new challenges are how to identify such a decision factor in OR and how to make a direct prediction for it in ML.

Following this idea, we investigate two crucial problems in marketing. The first one is the binary treatment assignment problem. When ignoring the cost incurred by the treatment (the cost-unaware version), the conditional average treatment effect (CATE) can be regarded as the decision factor. The common uplift models to predict CATE include meta-learners (Künzel et al. 2019; Nie and Wager 2021) and causal forests (Wager and Athey 2018; Athey, Tibshirani, and Wager 2019). The former are composed of multiple base models, and the latter usually combines generalized random forests (GRF) with double machine learning (DML) methods. Different from them, we propose a novel uplift model to make a direct prediction based on neural networks, which achieves good performance in both theory and practice. Despite the incremental revenue, the treatment can also incur different costs. In this cost-aware version, ROI (Return on Investment) of individuals can be regarded as the decision factor, which is calculated by the division of the incremental revenue and the incremental cost. However, most of existing works in causal inference did not involve the treatment cost and cannot apply to such a direct prediction. Although some works (Du, Lee, and Ghaffarizadeh 2019) investigated a similar problem to this, their loss function cannot converge to a stable extreme point in theory. In this paper, we design a convex loss function to guarantee an unbiased estimation of ROI of individuals when the loss converges.

As the second case study, we apply our approach to the budget allocation problem with multiple treatments and propose a novel evaluation metric for this problem in this paper. The Lagrange duality is an effective algorithm to solve the budget allocation problem. However, the decision factor in this algorithm contains the Lagrange multiplier that is uncertain and varies much with different budgets. The direct prediction for such a decision factor with all possible Lagrange multipliers is difficult and unrealistic. In this paper, we propose an equivalent algorithm to the Lagrange dual method while the decision factor in this algorithm is determined and irrelevant to the Lagrange multiplier. In addition, the corresponding causal learning model is developed and

a direct prediction for the decision factor can be obtained when the customized loss function converges. Finally, we also propose a novel evaluation metric named MT-AUCC to estimate the prediction result, which is similar to Area Under Uplift Curve (AUUC) (Rzepakowski and Jaroszewicz 2010) but involves both multiple treatments and incremental cost.

Large-scale simulations and online A/B Tests validate the effectiveness of our approaches. In the offline simulations, we use two real-world datasets collected from random control trials (RCT) in the online advertising/food delivery platforms. Multiple evaluation metrics and online AB tests show that our models and algorithms achieve significant improvement and increase the target reward by over 10% on average compared with state-of-the-art.

## Related Work

**Two-phase methods.** The composition of machine learning (ML) and operation research (OR) is one of the most common approaches to solve the resource allocation problem, which is called two-phase methods in this paper. In the first stage, uplift models are designed to predict the incremental response of individuals under different treatments. Besides meta-learners (Künzel et al. 2019; Nie and Wager 2021) and causal forests (Wager and Athey 2018; Athey, Tibshirani, and Wager 2019; Zhao, Fang, and Simchi-Levi 2017; Ai et al. 2022), representation learning (Johansson, Shalit, and Sontag 2016; Shalit, Johansson, and Sontag 2017; Yao et al. 2018) was also developed for uplift modeling. Instead of deriving an unbiased estimator, some works (Betlei, Diemert, and Amini 2021; Kuusisto et al. 2014) proposed a unified framework for learning to rank CATE. As one of the most effective algorithms, the Lagrangian duality is frequently used to solve many decision-making problems of different areas in the second stage. For example, it was developed to solve the budget allocation problem in marketing (Du, Lee, and Ghaffarizadeh 2019; Ai et al. 2022; Zhao et al. 2019) and compute the optimal bidding policy in online advertising (Hao et al. 2020).

**Direct learning methods.** Policy learning and reinforcement learning are two important methods to directly learn a treatment assignment policy rather than the treatment effect, which avoids the combination of ML and OR. A general framework for policy learning with observational data was proposed based on the doubly robust estimator (Athey and Wager 2021) and their work was extended to multi-action policy learning (Zhou, Athey, and Wager 2022). As a real-world application, the works (Xiao et al. 2019; Zhang et al. 2021) formulated the coupon allocation problem in sequential incentive marketing as a constrained Markov decision process and proposed reinforcement learning to solve it. However, all of the above methods moved the resource constraints into the reward function by using the Lagrangian multiplier. Hence, the model may need to be changed constantly with the variation of the Lagrangian multiplier.

**Decision-focused learning (DFL).** Similar to our motivation, DFL devotes itself to learning the model parameters based on the downstream optimization task rather than the prediction accuracy. Nevertheless, many existing works in

DFL required that the feasible region of the decision variables is fixed and known with certainty (Wilder, Dilkina, and Tambe 2019; Elmachtoub and Grigas 2022; Shah et al. 2022; Mandi et al. 2022). The most related work to ours is perhaps by Donti, Amos, and Kolter 2017, which addressed a stochastic optimization problem that contains both probabilistic and deterministic constraints. However, this work supposed that the decision variables are continuous and dealt with the probabilistic constraints by the Lagrangian duality, which is markedly distinct from ours.

### Binary Treatment Assignment Problem

We begin with a common marketing scenario, where part of  $M$  individuals are selected to receive the marketing action. We adopt the potential outcome framework (Sekhon 2008) to formulate this problem. Let  $X \in \mathbb{R}^d$  denote the feature vector and  $x$  its realization. Despite the incremental revenue, marketing actions can also incur significant costs. Let  $Y^r, Y^c$  denote the revenue and the cost respectively, and  $y^r, y^c$  its realization. Denote the treatment by  $T \in \{0, 1\}$  and its realization by  $t_i$ . Let  $(Y^r(1), Y^r(0))$  and  $(Y^c(1), Y^c(0))$  be the corresponding potential outcome when the individual receives the treatment or not. Define  $\tau^r(x_i), \tau^c(x_i)$  as the conditional average treatment effect, which can be calculated by

$$\tau^*(x_i) = E[Y^*(1) - Y^*(0)|X = x_i], * \in \{r, c\}.$$

Since most marketing actions have a positive effect on the response of an individual, we have  $\tau^r(x_i) > 0$  and  $\tau^c(x_i) > 0$ . The binary treatment assignment problem (BTAP) is to assign the treatment to part of the individuals to maximize the overall revenue on the platform, but requires that the incremental cost does not exceed a limited budget  $B$ . Let  $z_i \in \{0, 1\}$  be the decision variables. Therefore, BTAP can be formulated as the integer programming problem (1).

$$\begin{aligned} \max \quad & \sum_i z_i \tau^r(x_i) \\ \text{s.t.} \quad & \sum_i z_i \tau^c(x_i) \leq B \\ & z_i \in \{0, 1\}, \forall i. \end{aligned} \quad (1)$$

This problem is an equivalent 0/1 knapsack problem, which is NP-Hard. Fortunately, the simple Greedy algorithm (Algorithm 1) can achieve excellent performance whose approximation ratio satisfies  $\rho \geq 1 - \frac{\max_i \tau(x_i)}{\text{OPT}}$ , where OPT is the optimal solution of the problem (1).

**Definition 1.** *The decision factor of a combinatorial optimization algorithm is defined as a factor on which the final solution can be obtained by performing only the sorting or comparison operations.*

The decision factor is directly to the final solution of an algorithm and will be regarded as the learning objective in this paper. As is shown by Algorithm 1, the factor  $\tau^r(x_i)/\tau^c(x_i)$  can be taken as the decision factor, which is called as ROI (Return on Investment) of individual  $i$ .

---

### Algorithm 1: Greedy Algorithm for BTAP

---

**Input:**  $(\tau^r(x), \tau^c(x), B, M)$

**Output:** the solution to BTAP

- 1: Sort  $M$  individuals in descending order based on  $\frac{\tau^r(x_i)}{\tau^c(x_i)}$ .
  - 2: Assign the treatment to individual  $i$  in the sorted list until the budget  $B$  will be exceeded.
- 

### Cost-unaware Treatment Assignment Problem

When the treatment cost is nonexistent or the same for all the individuals (e.g., a push message), the prediction for ROI of individuals is reduced to the estimation of  $\tau^r(x_i)$ . The latter is an important problem in causal inference, and existing studies mainly involve meta-learners and causal forests. In this paper, we propose a novel uplift model to make a direct prediction for  $\tau^r(x_i)$  or the rank of  $\tau^r(x_i)$ , which can achieve good performance in both theory and practice.

Following the above notations, suppose that there is a data set of size  $N$  collected from random control trials (RCT) and denote the  $i$ -th sample by  $(x_i, t_i, y_i^r)$ . Denote the sample size that receive the treatment or not by  $N_1$  and  $N_0$ , respectively. Let  $s_i = \hat{h}(x_i)$  represent a score to rank  $\tau^r(x_i)$ , where  $\hat{h}(x_i)$  can be any machine learning model (e.g., linear regression or neural networks). Minimizing the loss function (2), we can obtain an unbiased estimation of  $\tau^r(x_i)$ . Due to the space limit, the detailed analysis is presented by Theorem 1 and its proof in Appendix A (Zhou et al. 2022).

$$\begin{aligned} \min L(s) = & -\left(\frac{1}{N_1} \sum_{i|t_i=1} y_i^r \ln q_i - \frac{1}{N_0} \sum_{i|t_i=0} y_i^r \ln q_i\right) \quad (2) \\ \text{s.t.} \quad & q_i = \frac{e^{s_i}}{\sum_i e^{s_i}}, \forall i \\ & s_i = \hat{h}(x_i) \in \mathbb{R}, \forall i \end{aligned}$$

**Theorem 1.** *When the loss function (2) converges,  $s_i$  can be used to rank  $\tau^r(x_i)$  and  $q_i = \frac{\tau^r(x_i)}{\sum_i \tau^r(x_i)}$  can be used to obtain an unbiased estimation of  $\tau^r(x_i)$ .*

### Cost-aware Treatment Assignment Problem

ROI of individuals is a composite object and most existing works predict it by the combination of multiple models. The latter may cause an enlargement of model errors due to the mathematical operations during combination. Therefore, we propose a novel learning model for direct ROI prediction.

$$\begin{aligned} \min L(s) = & -\left[\frac{1}{N_1} \sum_{i|t_i=1} (y_i^r \ln \frac{q_i}{1-q_i} + y_i^c \ln(1-q_i)) - \right. \\ & \left. \frac{1}{N_0} \sum_{i|t_i=0} (y_i^r \ln \frac{q_i}{1-q_i} + y_i^c \ln(1-q_i))\right] \quad (3) \\ \text{s.t.} \quad & q_i = \sigma(s_i), \forall i \\ & s_i = \hat{h}(x_i) \in \mathbb{R}, \forall i. \end{aligned}$$

Similarly, suppose that the data set  $(x_i, t_i, y_i^r, y_i^c)$  of size  $N$  is collected from RCT, and the count of the samples receiving the treatment or not is  $N_1$  and  $N_0$ , respectively. The

division operation may result in the high variance of ROI especially when  $\tau^c(x_i)$  is small. Therefore, the range of ROI is limited to  $(0, 1)$  by the scaling and truncating operations for  $Y^r$  or  $Y^c$  to decrease the risk of overfitting. Let  $s_i = \hat{h}(x_i)$  represent a score to rank ROI, where  $\hat{h}(x_i)$  can be any machine learning model. Define  $\sigma(\cdot)$  as the sigmoid function. The loss function (3) is designed to obtain an unbiased estimation of ROI or the rank of ROI for each individual. The detailed proof can be found in Theorem 2 and Appendix B (Zhou et al. 2022).

**Theorem 2.** *The loss function (3) is convex, and  $s_i$  can be used to rank ROI and  $q_i = \frac{\tau^r(x_i)}{\tau^c(x_i)}$  is an unbiased estimation of ROI of individual  $i$  when the loss converges.*

## Budget Allocation Problem with Multiple Treatments

In this section, we will discuss a more general marketing scenario, which was also introduced in many existing works. Most of them solve this problem by two-phase methods. Different from them, the machine learning model will be designed based on the decision factor to guarantee the consistency between the predictive objective and the optimization objective in this paper.

Suppose that there are multiples treatments and denote it by  $T \in \{1, 2, \dots, L\}$ . Let  $r_{ij}$  and  $c_{ij}$  be the revenue and the cost of individual  $i$  receiving treatment  $j$ , respectively. In marketing campaigns, multiple treatments usually refers to the different level of the treatment, e.g., the different discount of some products [1], the different count of gold pieces on the online video platform [2], the different price of cinema tickets in different regions [3] and so on. Suppose that  $j \in T$  represent the level of the marketing treatment and the treatment effect is larger if the level is higher. Therefore, we have  $r_{ij} < r_{ik}$  and  $c_{ij} < c_{ik}$  if  $j < k$  holds. Given a limited budget  $B$ , the budget allocation problem with multiple treatments (MTBAP) is to assign a certain treatment to each individual with the objective of optimizing the overall revenue on the platform. Let  $z_{ij} \in \{0, 1\}$  be the decision variable to denote whether to assign treatment  $j$  to individual  $i$ . Therefore, MTBAP can be formulated as the integer programming (4).

$$\begin{aligned} \max \quad & \sum_{ij} z_{ij} r_{ij} \\ \text{s.t.} \quad & \sum_{ij} z_{ij} c_{ij} \leq B \\ & \sum_j z_{ij} = 1, \forall i \\ & z_{ij} \in \{0, 1\}, \forall i, j. \end{aligned} \quad (4)$$

The first constraint refers to the limitation of the whole budget and the second requires that only one treatment can be assigned to each individual.

## Combinatorial Optimization Algorithm

This problem is a classical multiple choice knapsack problem and remains NP-Hard. Existing studies usually solve

this problem by using Lagrangian duality theory. Specifically, the upper bound of the original problem (4) can be obtained by solving the following dual problem (5).

$$\begin{aligned} \min_{\alpha \geq 0} \quad & \left( \begin{array}{l} \max \alpha B + \sum_{ij} z_{ij} (r_{ij} - \alpha c_{ij}) \\ \text{s.t.} \quad \sum_j z_{ij} = 1, \forall j \\ z_{ij} \in \{0, 1\}, \forall i, j \end{array} \right) \\ = \min_{\alpha \geq 0} \quad & \alpha B + \sum_i \max_j (r_{ij} - \alpha c_{ij}) \end{aligned} \quad (5)$$

The optimal  $\alpha^*$  for the dual problem can be derived by using the decreasing gradient algorithm or a binary search for  $\alpha$  with the terminal condition of  $B - \sum_{ij} z_{ij} c_{ij} \leq \epsilon$ . Given the optimal  $\alpha^*$ , an approximation solution for the original problem (4) is

$$\forall i, j, z_{ij} = 1 \iff j = \arg \max_j r_{ij} - \alpha^* c_{ij}. \quad (6)$$

According to Definition 1, the factor  $r_{ij} - \alpha^* c_{ij}$  is a decision factor for the Lagrangian duality algorithm and can be taken as the learning objective. However, the value of  $\alpha^*$  depends on the budget  $B$ , which is undetermined and varies much with the marketing environment. Hence, there is only a limited number of  $\alpha^*$  in the training dataset and it is possible that the optimal  $\alpha^*$  used in future campaigns has never been seen by the predictive model before, which dramatically decreases the precision of the model. Therefore, the direct prediction of the decision factor for all possible  $\alpha^*$  is difficult and unrealistic, and will not be considered in this paper.

However, another equivalent optimization algorithm can be derived from the above Lagrangian duality method. Before the details, we present an important assumption in economics, which is called as the Law of Diminishing Marginal Utility (Polleit 2011).

**Assumption 1** (The Law of Diminishing Marginal Utility). *The marginal utility of individuals decreases with the increasing investment of the marketing cost. Specifically, denote the marginal utility by*

$$l_{ij} = \frac{r_{ij+1} - r_{ij}}{c_{ij+1} - c_{ij}},$$

and we have  $l_{ij} \leq l_{ij-1}, \forall i, j$ .

Let  $(c_{ij}, r_{ij})$  be a point in rectangular coordinates. The value of  $r_{ij} - \alpha c_{ij}$  can be regarded as the projection of  $(c_{ij}, r_{ij})$  on  $(-\alpha, 1)$ . Let  $j^*$  be the solution of Lagrangian duality methods in Eq. (6), i.e.,  $j^* = \arg \max_j r_{ij} - \alpha c_{ij}$ . Based on Assumption 1, we can prove that it satisfies

$$l_{ij^*} \leq \alpha \leq l_{ij^*-1}.$$

Therefore, an equivalent optimization algorithm is obtained in Algorithm 2. Due to the space limit, the formal proof of Theorem 3 can be found in Appendix C (Zhou et al. 2022).

**Theorem 3.** *Algorithm 2 is equivalent to the Lagrangian duality method.*

Notice that the factor  $l_{ij}$  in Algorithm 2 can be taken as a decision factor, which is irrelevant to the Lagrangian multiplier  $\alpha$ . Therefore, the value of  $l_{ij}$  (or implicitly, the rank of  $l_{ij}$  for all  $i, j$ ) can be taken as the learning objective, which avoids the difficulty of directly predicting  $r_{ij} - \alpha c_{ij}$  for all possible  $\alpha$  in the Lagrangian duality method.

---

**Algorithm 2: An Equivalent Algorithm for MTBAP**


---

**Input:**  $(\ell_{ij}, B)$ 
**Output:** the solution to MTBAP

```

1:  $\alpha_l = 0, \alpha_r = \max_{ij} \ell_{ij}, z_{ij} = 0$  for  $\forall i, j$ .
2: /*The same binary search as Lagrangian duality.*/
3: while  $B - \sum_{ij} z_{ij} c_{ij} > \epsilon$  do
4:    $z_{ij} = 0$  for  $\forall i, j$ .
5:    $\alpha = (\alpha_l + \alpha_r)/2$ .
6:   for  $\forall i$  do {Equivalent to  $j^* = \arg \max_j r_{ij} - \alpha c_{ij}$ }
7:      $\ell_{i0} = +\infty, \ell_{iL} = 0$ .
8:     find  $j^*$  satisfying  $\ell_{ij^*} \leq \alpha < \ell_{ij^*-1}$ .
9:      $z_{ij^*} = 1$ .
10:  if  $B - \sum_{ij} z_{ij} c_{ij} >= 0$  then
11:     $\alpha_r = \alpha$ .
12:  else
13:     $\alpha_l = \alpha$ .
14: return  $z_{ij}$ .
```

---

**Direct Prediction Model**

Based on the above analysis, we propose a novel method for a direct prediction of  $\ell_{ij}$  in this subsection. Suppose that there is a data set  $(x_i, t_i, y_i^c, y_i^r)$  of size  $N$  collected from RCT, and denote by  $N_j$  the count of samples receiving treatment  $j$ . Similarly, the range of  $\ell_{ij}$  is limited to  $(0, 1)$  by scaling and truncating operations for  $Y^r$  or  $Y^c$  to reduce the risk of overfitting. Let  $s_{ij} = \tilde{h}(x_i, j)$  be the prediction of the rank of  $\ell_{ij}$ , where  $\tilde{h}(\cdot)$  is any machine learning model. Hence, minimize the loss function (7) and we can get the unbiased estimation of  $\ell_{ij}$ . The detailed analysis is shown by Theorem 4 and its proof in Appendix D (Zhou et al. 2022).

$$\begin{aligned}
\min L(s) = & -\left[ \sum_{j>1} \sum_{i|t_i=j} \frac{1}{N_j} (q_{ij-1} y_i^r - q_{ij-1}^2 y_i^c) - \right. \\
& \left. \sum_{j<L} \sum_{i|t_i=j} \frac{1}{N_j} (q_{ij} y_i^r - q_{ij}^2 y_i^c) \right] \quad (7) \\
\text{s.t.} \quad & q_{ij} = \sigma(s_{ij}), \forall i, j \\
& s_{ij} = \tilde{h}(x_i, j) \in \mathbb{R}, \forall i, j.
\end{aligned}$$

**Theorem 4.** When the loss function (7) converges,  $s_{ij}$  can be used to rank  $\ell_{ij}$  and  $q_{ij}$  can be used to obtain the unbiased estimation of  $\ell_{ij}$ .

**An Evaluation Metric**

Although AUUC (Area Under Uplift Curve) and AUCC (Area Under Cost Curve) (Du, Lee, and Ghaffarizadeh 2019) have been developed to evaluate the ranking performance of uplift models without/with treatment cost respectively, there is no evaluation metric for the estimation of marginal utilities ( $\ell_{ij}$ ) under different treatments. The latter is directly related to the business objective of MTBAP. Therefore, we propose a novel evaluation metric for such a purpose, which is called as MT-AUCC (Area Under Cost Curve for Multiple Treatments) in this paper.

**Cost Curve for Multiple Treatments.** Suppose that there is a model  $\mathcal{M} = f(x_i, j)$  to predict the value (or the rank) of

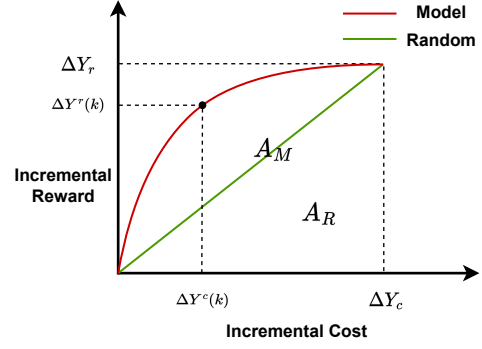


Figure 2: MT-AUCC

$\ell_{ij}$ . Firstly, we obtain two new quintuple sets based on  $\mathcal{M}$ .

- For each sample  $(x_i, t_i, y_i^c, y_i^r)$  with  $t_i < L$ , use model  $\mathcal{M}$  to obtain a score  $S_i = f(x_i, t_i)$  and get a new quintuple set  $\tilde{T} = \{(x_i, t_i, \alpha_i y_i^c, \alpha_i y_i^r, S_i) | t_i < L, \alpha_i = \frac{N}{N_{t_i}}\}$ ;
- for each sample  $(x_i, t_i, y_i^c, y_i^r)$  with  $t_i > 1$ , use this mode again to obtain  $S_i = f(x_i, t_i - 1)$  and get a quintuple set  $\tilde{C} = \{(x_i, t_i, \alpha_i y_i^c, \alpha_i y_i^r, S_i) | t_i > 1, \alpha_i = \frac{N}{N_{t_i}}\}$ .

Notice that the weight  $\alpha_i$  is used to balance the count of samples under different treatments. Next regard  $\tilde{T}$  and  $\tilde{C}$  as the new treatment group and control group respectively. Sort all the quintuples in  $\tilde{T}$  and  $\tilde{C}$  in descending order of  $S_i$ . For the top  $k$  quintuples in this sorted list, denote by  $\tilde{T}(k)$  (or  $\tilde{C}(k)$ ) the quintuples that belong to  $\tilde{T}$  (or  $\tilde{C}$ ). Therefore, the incremental cost  $\Delta Y^c(k)$  and reward  $\Delta Y^r(k)$  can be calculated for each point  $1 \leq k \leq |\tilde{T}| + |\tilde{C}|$  by Eq. (8).

$$\Delta Y^*(k) = \frac{k}{|\tilde{T}| + |\tilde{C}|} \left( \frac{\sum_{i \in \tilde{T}(k)} \alpha_i y_i^r}{|\tilde{T}(k)|} - \frac{\sum_{i \in \tilde{C}(k)} \alpha_i y_i^c}{|\tilde{C}(k)|} \right), \quad (8)$$

where  $* \in \{r, c\}$

As is shown in Fig. 2, take the tuple  $(\Delta Y^c(k), \Delta Y^r(k))$  as the coordinates and we can get a cost curve.

Denote by  $\Delta Y_c$  and  $\Delta Y_r$  the average incremental cost and reward of all the samples in  $\tilde{T}$  and  $\tilde{C}$ , which satisfies  $\Delta Y_* = \Delta Y^*(|\tilde{T}| + |\tilde{C}|)$ . For convenience of calculations, we can also use some points with  $p$  percent of  $\Delta Y_c$  to draw the cost curve. In addition, the X and Y axis in this curve can be normalized by being divided by  $\Delta Y_c$  and  $\Delta Y_r$  respectively.

**Area Under Cost Curve for Multiple Treatments (MT-AUCC).** Similar to AUUC and AUCC, the area under this cost curve can be regarded as an evaluation metric. As is shown by Fig. 2, denote by  $A_M$  and  $A_R$  the area under a model curve and a random benchmark curve respectively. In order to bound the result within  $(0, 1]$ , MT-AUCC of this model is defined as  $A_M/2A_R$ .

**Evaluation**

In this section, we will conduct large-scale offline and online numerical experiments to validate the performance of our models and algorithms.

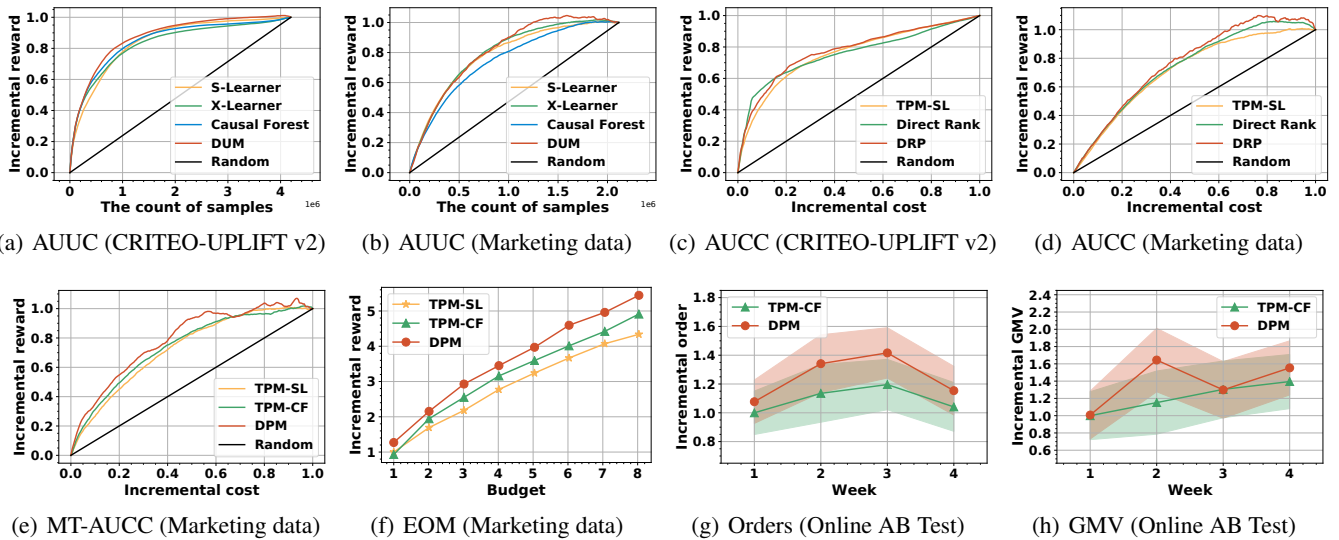


Figure 3: Evaluation Results

## Offline Simulation

**Dataset.** Two types of datasets are provided in this paper: an open real-world dataset and a marketing dataset collected from an online food delivery platform.

- CRITEO-UPLIFT v2. This dataset is provided by the AdTech company Criteo in the AdKDD’18 workshop (Diemert Eustache, Betlei Artem, Renaudin, and Massih-Reza 2018). The data is collected from a random control trial (RCT) that prevents a random part of users from being targeted by advertising. It contains 12 features, 1 binary treatment indicator and 2 response labels (visit/conversion). The (incremental) visit is regarded as the predictive objective for (cost-unaware) uplift modeling. In order to compare the performance of different models to predict ROI of individuals, we take the visit label as the cost and the conversion label as the reward. The whole dataset contains 13.9 million samples, and is randomly partitioned into two parts for train (70%) and test (30%), respectively.
- Marketing data. Money Off is a common marketing campaign in Meituan, an online food delivery platform. We conduct a four-week RCT in this platform where online shops will offer a random discount every day. Notice that the discount of a shop is the same for all the users to prevent price discrimination but may randomly change in different days, and different shops may offer different discounts. The data in the first two weeks is used for train and the others for test. The discount  $T \in \{0, 1, \dots, 6\}$  is taken as the treatment, where  $T = t$  means  $\$t$  cash off for each order whose price meets a given threshold. This dataset contains 75 features, 1 treatment label and 2 response labels (daily cost/orders). For the binary treatment assignment problem, we take the samples with  $T = 0$  as the control group and the samples with  $T > 0$  as the treatment group. For the budget allocation problem with multiple treatments, the budget refers to the whole cost

of all the shops and different discounts represent different treatments. This dataset contains 4.1 million samples.

**Evaluation Metric.** Multiple evaluation metrics are provided for offline evaluation in this experiment.

- AUUC (Area under Uplift Curve). A common metric to evaluate uplift models (Rzepakowski and Jaroszewicz 2010). In this experiment, the auuc score is computed by using CausalML packages (Chen et al. 2020).
- AUCC (Area under Cost Curve). A similar metric to AUUC, but designed for evaluating the performance to rank ROI of individuals (Du, Lee, and Ghaffarizadeh 2019).
- MT-AUCC. It is proposed in this paper and used to evaluate the performance of models to rank marginal utilities of different individuals under different treatments.
- EOM (Expected Outcome Metric). Based on RCT data, the expected outcome (reward/cost) can be obtained for arbitrary policy by using the computing methods in (Ai et al. 2022; Zhao, Fang, and Simchi-Levi 2017).

**Benchmark.** For each problem considered in this paper, multiple different models/algorithms are implemented and taken as the benchmarks.

- Cost-unaware binary treatment assignment problem
  - S-Learner. A single model predicting the response of individuals with/without the treatment. The CATE is computed by  $E(Y|X, T = 1) - E(Y|X, T = 0)$ .
  - X-Learner. A meta-learner approach proposed in (Künzel et al. 2019).
  - Causal Forest. An uplift model proposed in (Athey, Tibshirani, and Wager 2019). It is implemented here based on EconML packages (Battocchi et al. 2019).
  - DUM. The direct uplift modeling method in this paper.
- Cost-aware binary treatment assignment problem

- TPM-SL. The two-phase method which uses two S-Learner models to predict the incremental revenue and cost, respectively. Predict ROI of individuals by computing the ratio between these two models.
- Direct Rank. Similar to our work, a loss function is designed for ranking ROI of individuals in this model (Du, Lee, and Ghaffarizadeh 2019). However, we prove that it cannot achieve the correct rank when the loss converges in Appendix E (Zhou et al. 2022).
- DRP. The direct ROI prediction model in this paper.
- Budget allocation problem with multiple treatments
  - TPM-SL. The two-phase method mentioned in many existing works (Ai et al. 2022; Zhao et al. 2019). In the first stage, we use a S-Learner model to predict the response (reward/cost) of individuals under different treatments. In the second stage, the Lagrangian duality algorithm is developed to compute the approximately optimal solution.
  - TPM-CF. Instead of S-Learner, we use Causal Forests to predict the incremental response. It is implemented based on generalized random forests (GRF) in EconML packages (Battocchi et al. 2019), which can also support multiple treatments.
  - DPM. The approach in this paper that combines the direct prediction of marginal utilities and Algorithm 2.

The hyperparameters in these algorithms are obtained based on grid search and each data point in the experimental results is computed by running the programs for 20 times.

**Experimental Results.** For the cost-unaware binary treatment assignment problem, Fig. 3(a)-3(b) presents the comparison of four uplift models. First of all, our model DUM performs best in both CRITEO-UPLIFT v2 and Marketing data. As the common baseline, the result of S-Learner is not too bad. It is near to our algorithm DUM in both two datasets. For comparison, X-Learner and Causal Forest is not always superior to S-Learner. The former is worse in CRITEO-UPLIFT v2 and the latter is inferior to S-Learner in Marketing data. The detailed results can be found in Table 1 in Appendix F (Zhou et al. 2022).

Due to the robustness of S-Learner, it is still taken as the base model to predict ROI of individuals. As is shown by Fig. 3(c)-3(d), TPM-SL cannot perform well especially in Marketing data at this time. The incorrect loss function of Direct Rank causes that it cannot converge to a stable extreme point and is inferior to our model DRP. In Appendix E (Zhou et al. 2022), we will present the detailed analysis for the convergence of Direct Rank. Compared with TPM-SL and Direct Rank, our model DRP always performs best and achieves significant improvement.

Fig. 3(e)-3(f) shows the results of different models and algorithms to solve the budget allocation problem with multiple treatments. Since the tree-based uplift models were often used in many existing works (Ai et al. 2022; Zhao, Fang, and Simchi-Levi 2017) to deal with this problem, we also take TPM-CF as the baseline. Our approach DPM significantly outperforms TPM-SL and TPM-CF in MT-AUCC, which indicates that DPM is better at ranking marginal utilities. We

also use EOM to test the incremental reward of different approaches when given different budget in Fig. 3(f). In order to protect the data privacy of this platform, the budget and reward have been normalized. In spite of this, it is still clear that our approach DPM can always help the platform to obtain much more reward under different budget.

As is shown by Table 1 in Appendix F (Zhou et al. 2022), all the models proposed in this paper are more stable and have lower variance than other existing works. This is because our models can make a direct prediction for the final objective, and always converge to a stable extreme point.

## Online A/B Test

**Setup.** We deploy our algorithm (DPM) to support the Money Off campaign in Meituan (a food delivery platform), and conduct an online AB test for four weeks. There are 310k total shops in this experiment and they are randomly partitioned into three groups, named G-DPM, G-TPM and G-Control respectively. The discount  $T \in \{0, 1, \dots, 6\}$  is taken as the treatment and assigning a shop with treatment  $T = t$  means \$ $t$  cash off for each order whose price meets a given threshold. Given a limited budget, the objective is to decide the discount every day for each shop so as to maximize the total number of orders and GMV (Gross Merchandise Value) in this platform. Algorithm DPM and TPM-CF are deployed in the experiment groups named G-DPM and G-TPM, respectively. The group G-Control is taken as the control group and does not offer any discount. These groups will be randomly broken up every week. Therefore, the AB experiment is repeated four times and the period of each time is one week.

**Results.** Fig. 3(e)-3(f) shows the incremental orders and GMV relative to G-Control in each week. To protect data privacy, all the data points have been normalized that are divided by the incremental orders or GMV of TPM-SL in the first week. The shadow area in Fig. 3(e)-3(f) represents the confidence interval with a confidence level of 0.95, which is calculated by student’s t-test. Compared with TPM-CF, our approach DPM always performs better in incremental orders and is not inferior to it in incremental GMV in each week. To sum up, DPM achieves a significant growth by 14.3% in incremental orders and 13.6% in GMV on average.

## Conclusion

In this paper, we proposed a novel approach for solving resource allocation problems based on the decision factor. Taking it as the learning objective can avoid alternative mathematical operations performed on the prediction results. This idea was applied to solve two crucial problems in marketing and presented great advantages both theoretically and practically. Large-scale offline simulations and online AB tests validated the effectiveness of our approach.

Our future work will focus on the application of this approach in more complex marketing scenarios. For example, multiple marketing campaigns may be conducted at the same time and interact with each other. Therefore, deriving the decision factor and conducting direct heterogeneous causal learning in this situation are more challenging.



## Acknowledgments

This work was partially supported by National Key R&D Program of China (2022YFB2901800), China NSF grants (62172206) and China NSF of Jiangsu Province (BK20201248).

## 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 *The ACM Web Conference (WWW)*, 2310–2319.
- 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.
- Battocchi, K.; Dillon, E.; Hei, M.; Lewis, G.; Oka, P.; Opreescu, M.; and Syrgkanis, V. 2019. EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/microsoft/EconML>. Accessed: 2023-03-10.
- Betlei, A.; Diemert, E.; and Amini, M.-R. 2021. Uplift Modeling with Generalization Guarantees. In *the ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD)*, 55–65.
- Chen, H.; Harinen, T.; Lee, J.-Y.; Yung, M.; and Zhao, Z. 2020. CausalML: Python Package for Causal Machine Learning. arXiv:2002.11631.
- Diemert Eustache, Betlei Artem; Renaudin, C.; and Massih-Reza, A. 2018. A Large Scale Benchmark for Uplift Modeling. In *the ACM AdKDD and TargetAd Workshop*. ACM.
- Donti, P.; Amos, B.; and Kolter, J. Z. 2017. Task-Based End-to-End Model Learning in Stochastic Optimization. *Advances in Neural Information Processing Systems (NIPS)*, 30.
- Du, S.; Lee, J.; and Ghaffarizadeh, F. 2019. Improve User Retention with Causal Learning. In *the ACM SIGKDD Workshop on Causal Discovery*, volume 104, 34–49. PMLR.
- Elmachtoub, A. N.; and Grigas, P. 2022. Smart “Predict, then Optimize”. *Management Science*, 68(1): 9–26.
- Hao, X.; Peng, Z.; Ma, Y.; Wang, G.; Jin, J.; Hao, J.; Chen, S.; Bai, R.; Xie, M.; Xu, M.; et al. 2020. Dynamic Knapsack Optimization towards Efficient Multi-Channel Sequential Advertising. In *International Conference on Machine Learning (ICML)*, 4060–4070. PMLR.
- Hua, J.; Yan, L.; Xu, H.; and Yang, C. 2021. Markdowns in E-Commerce Fresh Retail: A Counterfactual Prediction and Multi-Period Optimization Approach. In *The ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 3022–3031.
- Johansson, F.; Shalit, U.; and Sontag, D. 2016. Learning Representations for Counterfactual Inference. In *International Conference on Machine Learning (ICML)*, 3020–3029. PMLR.
- Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for Estimating Heterogeneous Treatment Effects using Machine Learning. *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 *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 50–65. Springer.
- Mandi, J.; Bucarey, V.; Tchomba, M. M. K.; and Guns, T. 2022. Decision-Focused Learning: Through the Lens of Learning to Rank. In *International Conference on Machine Learning (ICML)*, 14935–14947. PMLR.
- Nie, X.; and Wager, S. 2021. Quasi-oracle Estimation of Heterogeneous Treatment Effects. *Biometrika*, 108(2): 299–319.
- Polleit, T. 2011. What Can the Law of Diminishing Marginal Utility Teach Us. *Mises Institute*.
- Rzepakowski, P.; and Jaroszewicz, S. 2010. Decision Trees for Uplift Modeling. In *The IEEE International Conference on Data Mining (ICDM)*, 441–450. IEEE.
- Sekhon, J. S. 2008. The Neyman-Rubin Model of Causal Inference and Estimation via Matching Methods. *The Oxford Handbook of Political Methodology*, 2: 1–32.
- Shah, S.; Wang, K.; Wilder, B.; Perrault, A.; and Tambe, M. 2022. Decision-Focused Learning without Decision-Making: Learning Locally Optimized Decision Losses. In *Advances in Neural Information Processing Systems (NIPS)*.
- Shalit, U.; Johansson, F. D.; and Sontag, D. 2017. Estimating Individual Treatment Effect: Generalization Bounds and Algorithms. In *International Conference on Machine Learning (ICML)*, 3076–3085. PMLR.
- 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.
- Wilder, B.; Dilkina, B.; and Tambe, M. 2019. Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 1658–1665. AAAI Press.
- 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 *the ACM International Conference on Information and Knowledge Management (CIKM)*, 971–980.
- 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 (NIPS)*, 31.
- Zhang, Y.; Tang, B.; Yang, Q.; An, D.; Tang, H.; Xi, C.; LI, X.; and Xiong, F. 2021. BCORLE( $\lambda$ ): An Offline Reinforcement Learning and Evaluation Framework for Coupons Allocation in E-commerce Market. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 20410–20422.



Zhao, K.; Hua, J.; Yan, L.; Zhang, Q.; Xu, H.; and Yang, C. 2019. A Unified Framework for Marketing Budget Allocation. In *the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 1820–1830.

Zhao, Y.; Fang, X.; and Simchi-Levi, D. 2017. Uplift Modeling with Multiple Treatments and General Response Types. In *International Conference on Data Mining (ICDM)*, 588–596. SIAM.

Zhou, H.; Li, S.; Jiang, G.; Zheng, J.; and Wang, D. 2022. Direct Heterogeneous Causal Learning for Resource Allocation Problems in Marketing. *arXiv preprint*, arXiv:2211.15728.

Zhou, Z.; Athey, S.; and Wager, S. 2022. Offline Multi-Action Policy Learning: Generalization and Optimization. *Operations Research*.