

A Unified Framework for Marketing Budget Allocation

Kui Zhao, Junhao Hua, Ling Yan, Qi Zhang, Huan Xu, Cheng Yang

Machine Intelligence Technology Lab, Alibaba Group

Hangzhou, Zhejiang, China

{zhaokui.zk,junhao.hjh,yanling.yl,qiyue.zq,huan.xu,charis.yangc}@alibaba-inc.com

ABSTRACT

While marketing budget allocation has been studied for decades in traditional business, nowadays online business brings much more challenges due to the dynamic environment and complex decision-making process. In this paper, we present a novel unified framework for marketing budget allocation. By leveraging abundant data, the proposed data-driven approach can help us to overcome the challenges and make more informed decisions. In our approach, a semi-black-box model is built to forecast the dynamic market response and an efficient optimization method is proposed to solve the complex allocation task. First, the response in each market-segment is forecasted by exploring historical data through a semi-black-box model, where the capability of logit demand curve is enhanced by neural networks. The response model reveals relationship between sales and marketing cost. Based on the learned model, budget allocation is then formulated as an optimization problem, and we design efficient algorithms to solve it in both continuous and discrete settings. Several kinds of business constraints are supported in one unified optimization paradigm, including cost upper bound, profit lower bound, or ROI lower bound. The proposed framework is easy to implement and readily to handle large-scale problems. It has been successfully applied to many scenarios in Alibaba Group. The results of both offline experiments and online A/B testing demonstrate its effectiveness.

CCS CONCEPTS

• **Information systems** → **Decision support systems**; • **Applied computing** → *Marketing*; • **Computing methodologies** → Supervised learning; Planning and scheduling.

KEYWORDS

Marketing; Budget allocation; Market response; Forecasting; Optimization

1 INTRODUCTION

When and where should you spend the money? The answer to this question is the key to budget allocation in marketing and this topic has been studied for decades in traditional business. The environment and decision-making process in online business are much more dynamic and complex than traditional ones. And the marketing budget allocation is sophisticated and needs to be readjusted quite often, e.g. weekly or even daily, much more often than a few times per year like traditional marketing does. That brings great challenges, especially for companies with large and diversified markets [19, 32]. Fortunately, in online business, commercial activities can be tracked and depicted by the collected data from multiple sources. By leveraging these data, data-driven methods can help us to overcome above challenges. The intelligent techniques allow us

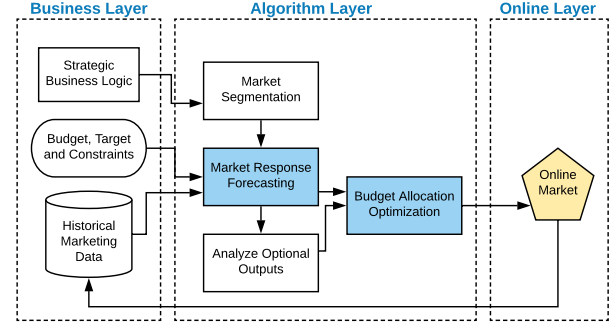


Figure 1: Our framework for marketing budget allocation.

to gain deeper insights into the market and make more informed decisions. Moreover, when the market is monitored and analyzed automatically, new patterns can be detected more quickly, and the budget allocation is thus readjusted timely [1]. In this paper, we present a novel unified framework for marketing budget allocation, which contains two sequential steps: learning market response models from historical data, and optimizing budget allocation based on learned models. The whole framework is illustrated in Figure 1.

In marketing, the whole market is divided into many segments, according to different commodities, customers, and consumption time, etc. In order to resolve how to allocate the budget into each market-segment, we need to forecast sales response of each segment if the budget is spent on it. Although black-box forecasting methods, such as neural networks, have been widely used in many applications [21, 33, 35], there are several gaps between forecasting and decision making [2]. One of the biggest challenges is that it is too difficult to translate black-box forecasting into allocation decisions. In contrast, the logit demand curve has an explicit expression and it is very popular in economics [26, 28, 36]. However, fitting logit demand curves in each market-segment separately has no capability of sharing information among different segments. That may cause serious data sparseness problem. Therefore, we propose a semi-black-box model to extend the capability of logit demand curves via neural networks. The input of forecasting is firstly split into two different parts: independent variables and contextual variables. The independent variable we have here is the budget allocated to each segment. The contextual variables are intrinsic attributes of market, such as brands of commodities, living cities of customers, consumption time and so on, and they are constants in each segment. It will be shown that the elasticity at market average cost (or market cost) of logit response model is determined by the bias parameter. By taking contextual variables as input, we use the same neural network to calculate bias parameters for all segments.

Based on that, fitting logit demand curves is extended to a process of learning the elasticity at market cost together with a specific sensitivity coefficient for each segment, where the relationship between sales and cost is depicted in an explicit way.

As long as response models are acquired, budget allocation can be formulated as an optimization problem. Let N denote the number of market-segments, budget allocation needs to optimize over N variables to maximize sales under some budget constraint. Due to the non-convexity of logit demand curves, the optimization problem is non-convex. Inspired by [7, 20], we reformulate the problem into an equivalent convex optimization problem. By introducing a dual variable, the Lagrange multiplier method is applied. According to KKT conditions, the original problem is reduced to a root finding problem with respect to dual variable, and it can be solved by bisection on one dual variable, instead of searching N primal variables. The algorithm converges in a few iterations and the complexity of each iteration is $O(N)$, thus it is readily to handle large-scale problems. Several kinds of business constraints, including cost upper bound, profit lower bound, or ROI lower bound, are supported in one unified paradigm with minor adaptations. What's more, the optimization paradigm can be easily extended to discrete settings with a little additional computation, where budget variables can only be chosen from a set of discrete values. One major cause of discrete decision variables is that 99-ending display prices, e.g. \$99, \$199 or \$299, are usually preferred after discount or premium [30] in many situations. We will show that the discrete setting can be solved by constructing a Multiple Choice Knapsack Problem (MCKP) [17] from the solution of relaxed continuous problem. Dyer et al. [9] and Zemel et al. [38] independently developed approximation algorithms for MCKP, running in $O(N)$ time, and the algorithm can be further boosted by pruning [27].

The proposed framework is easy to implement and has been successfully applied to many scenarios in Alibaba Group¹:

- *Daily Marketing*: Taopiaopiao, Koubei, CBU, Fliggy.
- *Shopping Festival*: 11.11 in Tmall and 12.12 in Lazada.
- *Interest Discount*: Credit Pay and Cash Now in Ant Financial.

To demonstrate the effectiveness of our framework, we present results of both offline experiments and online A/B testing. The offline part includes experiments on public datasets and optimization simulations. For the online part, we would like to take the daily marketing in Taopiaopiao as a typical example, on which an A/B testing is conducted for five weeks and the improvements are remarkable. In addition, it is worth mentioning that there are billions of decision variables in several scenarios, such as Figgy, Credit Pay, Cash Now, etc., and our approach can handle them very well.

The rest of our paper is organized as follows. We describe how to learn the market response in Section 2, and how to allocate budget in Section 3. We present the experimental setups and results in Section 4. We briefly review related work in Section 5. The conclusions and future plans are given in Section 6. Finally, details related to reproducibility are given in Supplement Section.

¹ Taopiaopiao is a platform for selling movie tickets. Koubei is a platform for local services. CBU (also called 1688) is a business-to-business (B2B) platform. Fliggy is a platform for selling travel tickets. Tmall is a business-to-consumer (B2C) platform, and Lazada is the Southeast Asian version. Credit Pay (also called Huabei) provides virtual credit card service, and Cash Now (also called Jiebei) provides consumer loan service.




Figure 2: The logit response model.

2 MARKET RESPONSE

In online business, the commercial activities can be tracked and depicted by collected data from multiple sources. Based on that data, the whole market can be divided into more granular segments than traditional ones. Each segment may consist of very specific commodities and customers. For instance, black wool sweater for children is a market-segment. The market can be further segmented according to consumption time, e.g. movie tickets on Saturday evening at a specific cinema is another example of market-segment.

2.1 Logit Response Model

The fundamental building block of marketing budget allocation is to forecast sales response of each segment if the budget is spend on it. As black-box forecasting methods, neural networks are very powerful and have wide applications [21, 33, 35]. However, there are many challenges to translate black-box forecasting into allocation decisions [2]. In contrast, the demand curves in economics usually have explicit expressions. There are many kinds of demand curves, such as linear, log-linear, constant-elasticity, logit and so on [26, 34]. In this work, we focus on logit demand curve, which is the most popular response model [26, 28, 36].

The logit model for i -th segment is defined as follow and one sample is illustrated in Figure 2:

$$d_i(c) := \frac{D_i}{1 + \exp\{-(a_i + b_i c)\}}, i = 1, \dots, N, \quad (1)$$

where c is the unit cost on marketing, and $D_i > 0$ indicates the overall size of i -th market-segment. c can be either greater than or less than 0. Positive cost on marketing means price discount and negative cost means price premium. Since the size of each market-segment is quite stable, D_i can be calculated by simple statistics in advance. So the only parameters need to be learned are a_i, b_i . Specifically, b_i indicates sensitivity and in normal situations we have $b_i > 0$. That is to say greater cost on marketing brings higher sales. The response curve is steepest at point $\hat{c}_i = -a_i/b_i$ (as indicated in Figure 2), and we call this point as *market cost*, which represents the average cost on marketing (including all competitors) in the whole i -th market-segment. The sales is very sensitive to cost when the cost is close to \hat{c}_i . In another word, a substantial shifts in sales could be expected even if the change in cost is very small. The sales becomes less sensitive to cost when cost goes away from the market cost, and this phenomenon is called as *diminishing marginal effect*.

The most common measure of sensitivity is elasticity. Inspired by *price elasticity* [26], we define *cost elasticity* as the ratio of negative percentage change in sales to percentage change in cost:

$$e_i(c_1, c_2) := -\frac{[d_i(c_2) - d_i(c_1)]c_1}{[c_2 - c_1]d_i(c_1)}. \quad (2)$$

The elasticity at c can be derived by taking the limit of Eq. (2) as c_2 approaches c_1 :

$$e_i(c) = -\nabla_c d_i(c) \frac{c}{d_i(c)}. \quad (3)$$

THEOREM 1. *The elasticity at market cost $\hat{c}_i = -a_i/b_i$ for logit response model is equal to $-a_i/2$.*²

2.2 Semi-black-box Model

The capability of fully utilizing data for logit model is limited. More specifically, there is no capability of sharing information across different segments when we learn a_i and b_i for each market-segment i separately. For market-segments with few distinct values of historical budget allocation, that may cause serious data sparseness problem. Actually, the input variables of forecasting model can be split into two different parts: independent variables and contextual variables. The independent variable we have here is c or the budget allocated to each segment. The contextual variables are intrinsic attributes of market and we denote it as \mathbf{x}_i . The contextual variables include brands of commodities, categories of commodities, living cities of customers, and consumption time, etc., and they are not changing with the independent variable in each segment. The neural networks take both \mathbf{x}_i and c as input, and connect them with sales in a black-box way. Contrarily, the logit model only takes c as input, and connect it with sales in an explicit way. We propose a semi-black-box model, which is able to take both \mathbf{x}_i and c as input, as well as keep the connection between c and sales in an explicit.

ASSUMPTION 1. *The elasticity at market cost for logit response model is determined by contextual variables.*

It is intuitive and reasonable to introduce the above assumption. For instance, T-shirts are less elastic than sweaters in summer, and movie tickets on the weekend is less elastic than that on the weekday. Instead of directly forecasting the final sales, we use a neural network $e(\mathbf{x}_i)$ to calculate the elasticity at market cost, i.e. $-a_i/2$, for each market-segment $i(i = 1, \dots, N)$. The elasticity information can be shared across different segments through $e(\mathbf{x}_i)$. Then the traditional logit response model is extended to a semi-black-box model:

$$d_i(c) = \frac{D_i}{1 + \exp\{2e(\mathbf{x}_i) - b_i c\}}, \quad (4)$$

which integrates the capability of neural networks and explicitness of logit demand curves.

The parameters in neural network $e(\mathbf{x}_i)$ as well as $b_i(i = 1, \dots, N)$ need to be learned. We first define *market share* $q_i(c)$ as follow:

$$q_i(c) := \frac{d_i(c)}{D_i} = \frac{1}{1 + \exp\{-(a_i + b_i c)\}}, \quad (5)$$

where $a_i = -2e(\mathbf{x}_i)$. As D_i can be calculated by simple statistics in advance, we can easily construct training set $\{([\mathbf{x}_i, c_1], \hat{q}_{i1}), \dots, ([\mathbf{x}_i, c_{M_i}], \hat{q}_{iM_i})\}$ for each market-segment $i(i = 1, \dots, N)$ from

²The proofs of all theorems can be found in the supplement.

historical observations. Then parameters are learned by minimizing the following negative log-likelihood through gradient decent:

$$\mathcal{J} = -\frac{1}{\sum_{i=1}^N M_i} \sum_{i=1}^N \sum_{j=1}^{M_i} [\hat{q}_{ij} \log q_{ij} + (1 - \hat{q}_{ij}) \log(1 - q_{ij})]. \quad (6)$$

3 BUDGET ALLOCATION

As long as response models are acquired, budget allocation can be formulated as an optimization problem.

3.1 Problem Formulation

Suppose there are N market-segments, budget allocation needs to decide unit cost c_i for each segment $i(i = 1, \dots, N)$ under some budget constraint. Let $\mathbf{c} := [c_1, \dots, c_N]$, our objective is to find an optimal \mathbf{c} maximizing total sales under budget constraint. Formally,

$$\begin{aligned} \min_{\mathbf{c}} & -\sum_{i=1}^N d_i(c_i), \\ \text{s.t.} & \sum_{i=1}^N d_i(c_i)c_i \leq B, \end{aligned} \quad (7)$$

where B is the budget bound and it can be either greater than or less than 0. Positive B means cost upper bound and negative B means profit lower bound.

Because $\sum_{i=1}^N d_i(c_i)c_i = \sum_{i=1}^N D_i q_i(c_i)c_i$ is non-convex with respect to c_i , directly solving the above problem is difficult. Inspired by [7, 20], we can reformulate the problem into an equivalent convex optimization problem. Since $q_i(c_i)$ obviously is a strictly increasing function when $b_i > 0$, we can get its inverse function, which maps market share q_i to marketing cost c_i , as follow:

$$\begin{aligned} c_i(q_i) &= -\frac{a_i}{b_i} - \frac{1}{b_i} [\ln(1 - q_i) - \ln q_i], \\ q_i &\in (0, 1), i = 1, \dots, N. \end{aligned} \quad (8)$$

Then we define $\mathbf{q} := [q_1, \dots, q_N]$ and constraint function $g(\mathbf{q})$ as:

$$g(\mathbf{q}) := \sum_{i=1}^N D_i q_i c_i(q_i) - B. \quad (9)$$

THEOREM 2. *$g(\mathbf{q})$ is strongly convex with respect to \mathbf{q} .*

Therefore, we can reformulate the problem as follow:

$$\begin{aligned} \min_{\mathbf{q}} & -\sum_{i=1}^N D_i q_i, \\ \text{s.t.} & g(\mathbf{q}) \leq 0, \end{aligned} \quad (10)$$

which is a convex optimization problem with respect to \mathbf{q} .

3.2 Algorithm

By introducing a dual variable λ , the Lagrange multiplier method is applied to solve the convex optimization problem in Eq. (10):

$$\begin{aligned} \mathcal{L}(\mathbf{q}, \lambda) &= -\sum_{i=1}^N D_i q_i + \lambda \left(\sum_{i=1}^N D_i q_i c_i(q_i) - B \right) \\ &= \sum_{i=1}^N (\lambda D_i q_i c_i(q_i) - D_i q_i) - \lambda B. \end{aligned} \quad (11)$$

3.2.1 *KKT conditions.* Both objective function and constraint function are differentiable and convex, and the KKT conditions are:

$$g(\mathbf{q}^*) \leq 0, \quad (12a)$$

$$\lambda^* \geq 0, \quad (12b)$$

$$\lambda^* g(\mathbf{q}^*) = 0, \quad (12c)$$

$$\nabla_{q_i}(\lambda^* D_i q_i^* c_i(q_i^*) - D_i q_i^*) = 0, i = 1, \dots, N, \quad (12d)$$

where \mathbf{q}^*, λ^* are the optimal solutions of primal and dual problems respectively, with zero duality gap [5].

While $g(\mathbf{q})$ is strongly convex, the minimum value of $g(\mathbf{q})$ exists, and let $g(\tilde{\mathbf{q}})$ denote it. To obtain the minimum value, we take the partial derivative of $g(\mathbf{q})$ with respect to each q_i and set it to 0. Then, we can get:

$$\ln \frac{\tilde{q}_i}{1 - \tilde{q}_i} + \frac{\tilde{q}_i}{1 - \tilde{q}_i} = a_i - 1, i = 1, \dots, N. \quad (13)$$

Using Lambert W function [6], $g(\tilde{\mathbf{q}})$ is *uniquely* attained at:

$$\tilde{q}_i = \frac{W(\exp(a_i - 1))}{W(\exp(a_i - 1)) + 1}, i = 1, \dots, N, \quad (14)$$

where $W(\cdot)$ is the Lambert W function. When $g(\tilde{\mathbf{q}}) > 0$, there is no solution. When $g(\tilde{\mathbf{q}}) = 0$, the optimal solution is imminently obtained at $\mathbf{q}^* = \tilde{\mathbf{q}}$. When $g(\tilde{\mathbf{q}}) < 0$, the following assumption holds:

ASSUMPTION 2. *There exists \mathbf{q} such that $g(\mathbf{q}) < 0$.*

When the above assumption holds, the constraint function $g(\mathbf{q})$ satisfies Slater's condition, and the KKT conditions in Eq. (12) provide necessary and sufficient conditions for optimality [5].

THEOREM 3. *The fourth KKT condition Eq. (12d) is equivalent to $\lambda^* \neq 0$ and $q_i^* = \frac{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1))}{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1)) + 1}$ ($i = 1, \dots, N$), where $W(\cdot)$ is the Lambert W function.*

From $\lambda^* \neq 0$ and Eq. (12c), it is easy to conclude that $g(\mathbf{q}^*) = 0$. Then we obtain the following optimality conditions:

$$g(\mathbf{q}^*) = \sum_{i=1}^N D_i q_i^* c_i(q_i^*) - B = 0, \quad (15a)$$

$$q_i^* = \frac{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1))}{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1)) + 1}, i = 1, \dots, N, \quad (15b)$$

$$\lambda^* > 0. \quad (15c)$$

We can see that the primal variables $q_i^* (i = 1, \dots, N)$ can be represented by the same dual variable λ^* . That is to say searching N primal variables is equal to searching one dual variable λ . Let us define:

$$q_i(\lambda) := \frac{W(\exp(a_i + \frac{b_i}{\lambda} - 1))}{W(\exp(a_i + \frac{b_i}{\lambda} - 1)) + 1}, \quad (16)$$

$$\lambda \in (0, +\infty), i = 1, \dots, N,$$

and

$$\begin{aligned} f(\lambda) &:= \sum_{i=1}^N D_i q_i(\lambda), \\ g(\lambda) &:= g(\mathbf{q}(\lambda)), \lambda \in (0, +\infty). \end{aligned} \quad (17)$$

Algorithm 1 Bisection algorithm for budget allocation

Input: $\epsilon, B, a_i, b_i, D_i, i = 1, \dots, N$.

```

1:  $\lambda_l = 0, \lambda_r = 1$ .
2: while  $g(\lambda_r) > 0$  do
3:    $\lambda_l = \lambda_r, \lambda_r = 2\lambda_r$ .
4: end while
5: repeat
6:    $\lambda_m = (\lambda_l + \lambda_r)/2$ .
7:   if  $g(\lambda_m) > 0$  then
8:      $\lambda_l = \lambda_m$ .
9:   else
10:     $\lambda_r = \lambda_m$ .
11:   end if
12: until  $\lambda_r - \lambda_l \leq \epsilon$ 
13:  $\lambda^* = \lambda_r, \mathbf{q}^* = \mathbf{q}(\lambda^*)$ .
14: Compute  $c_i^* = c_i(q_i^*)$  using Eq. (8).
```

If and only if the root λ^* of $g(\lambda) = 0$ is found, the problem is solved and the optimal solution is $\mathbf{q}^* = \mathbf{q}(\lambda^*)$.

3.2.2 *Bisection method.* The bisection method is a numerical root-finding method, and the precondition is to know two values with opposite signs. The process repeatedly bisects the interval defined by these two values and breaks when the interval is sufficiently small.

THEOREM 4. *The functions $f(\lambda), g(\lambda)$ are strictly decreasing with respect to λ , and $\lim_{\lambda \rightarrow 0} g(\lambda) > 0$.*

When Assumption 2 holds, there exists $\lambda < +\infty$ such that $g(\lambda) \leq 0$. Together with Theorem 4, the root λ^* of $g(\lambda) = 0$ can be found by bisection, which is shown in Algorithm 1, and one example is illustrated in Figure 3. As seen from above, we approximately set $\lambda^* = \lambda_r$ in order to satisfy the constraint. In practice, the algorithm can also be early terminated by $f(\lambda_l) - f(\lambda_r) \leq \epsilon'$, besides $\lambda_r - \lambda_l \leq \epsilon$. It is obvious that $f(\lambda^*) - f(\lambda_r) < \epsilon'$ when $f(\lambda_l) - f(\lambda_r) \leq \epsilon'$.

3.2.3 *Time complexity.* At line 2 – 4 of Algorithm 1, there are at most $\lceil \log_2(\max(\lambda^*, 1)) \rceil + 1$ iterations to find λ_r satisfying $g(\lambda_r) \leq 0$. At line 5 – 12, there are at most $\lceil \log_2(\frac{\max(\lambda^*, 1)}{\epsilon}) \rceil$ iterations. Therefore, the total number of iterations is less than $\lceil 2 \log_2(\max(\lambda^*, 1)) \rceil + \lceil \log_2(\frac{2}{\epsilon}) \rceil$. If the algorithm can be further early terminated by $f(\lambda_l) - f(\lambda_r) \leq \epsilon'$, there are at most $\lceil \log_2(\min(\frac{\max(\lambda^*, 1)}{\epsilon}, \frac{\sum_{i=1}^N D_i}{\epsilon'})) \rceil$ iterations at line 5 – 12. So the total number of iteration is less than a constant, i.e. $\lceil \log_2(\max(\frac{1}{\epsilon}, \frac{\sum_{i=1}^N D_i}{\epsilon'})) \rceil + 1$. In practice, the algorithm converges in a few iterations and the complexity of each iteration is $O(N)$, thus it is readily to handle large-scale problems.

3.3 Extension

3.3.1 *ROI.* In many situations, it is not easy to determine the specific B . Instead, the efficiency ratio of spending budget is expected to be no less than a given value. Let $R > 0$ denote the lower bound of efficiency ratio or Return on Investment (ROI), the original problem




Figure 3: The plots of objective and exceeded budget with respect to dual variable λ . The optimal objective value $f(\lambda^*)$ is attained when $g(\lambda^*) = 0$.

in Eq. (7) is extended to:

$$\begin{aligned} \min_{\mathbf{c}} \quad & \sum_{i=1}^N d_i(c_i), \\ \text{s.t.} \quad & \sum_{i=1}^N (Rd(c_i)c_i - d_i(c_i)) \leq 0. \end{aligned} \quad (18)$$

Then the constraint function in Eq. (9) can be replaced by:

$$g'(\mathbf{q}) := \sum_{i=1}^N (Rd_i q_i c_i(q_i) - D_i q_i). \quad (19)$$

THEOREM 5. $g'(\mathbf{q})$ is strongly convex with respect to \mathbf{q} . $g'(\lambda)$ is strictly decreasing with respect to λ , and $\lim_{\lambda \rightarrow 0} g'(\lambda) > 0$.

After a similar derivation, we can conclude that the new optimization problem still can be solved by Algorithm 1, except $q_i(\lambda)$ in Eq. (16) is replaced by:

$$q_i(\lambda) := \frac{W(\exp(a_i + \frac{b_i}{\lambda R} + \frac{b_i}{R} - 1))}{W(\exp(a_i + \frac{b_i}{\lambda R} + \frac{b_i}{R} - 1)) + 1}, \quad (20)$$

$\lambda \in (0, +\infty), i = 1, \dots, N,$

and the input B is no longer required.

3.3.2 Discrete settings. In many situations, each c_i can only be chosen from a set of discrete values. One common reason is that 99-ending display prices (e.g. \$99, \$199 or \$299) are preferred after discount or premium [30]. Let $S_i := \{c_i^1, \dots, c_i^{|S_i|}\}$ denote the set of optional values for c_i , the original problem in Eq. (7) is then extended to a NP-hard problem:

$$\begin{aligned} \min_{\mathbf{c}} \quad & \sum_{i=1}^N d_i(c_i), \\ \text{s.t.} \quad & \sum_{i=1}^N d_i(c_i)c_i \leq B, \\ & c_i \in S_i, i = 1, \dots, N. \end{aligned} \quad (21)$$

To solve the discrete problem, we first solve the relaxed continuous problem in Eq. (7) and let $\mathbf{c}^* = [c_1^*, \dots, c_N^*]$ denote its solution. After defining:

$$\begin{aligned} c_i^l &:= \arg \min_{c_i \in S_i, c_i \leq c_i^*} |c_i - c_i^*|, \\ c_i^u &:= \arg \min_{c_i \in S_i, c_i > c_i^*} |c_i - c_i^*|, i = 1, \dots, N. \end{aligned} \quad (22)$$

The discrete problem can be then solved by constructing a Multiple Choice Knapsack Problem (MCKP) [17], where B is the capacity of knapsack. In the constructed MCKP, i -th class represents i -th market-segment, and each class contains two items: the profit and weight of the first item are $d_i(c_i^l), d_i(c_i^l)c_i^l$; the profit and weight of the second item are $d_i(c_i^u), d_i(c_i^u)c_i^u$. Since $\sum_{i=1}^N d_i(c_i^l)c_i^l \leq \sum_{i=1}^N d_i(c_i^*)c_i^* \leq B$, there exists a solution to the constructed MCKP. Dyer et al. [9] and Zemel et al. [9] independently developed approximation algorithms (called Dyer-Zemel) for MCKP, running in $O(N)$ time, and the algorithm can be further boosted by pruning [27]. So the additional computation is very limited. In practice, when $|S_i|$ is small we suggest to directly solve the discrete setting by Dyer-Zemel algorithm with multiple items in each class, instead of solving the relaxed continuous problem first. Actually, when discrete setting is directly solved by Dyer-Zemel algorithm, profit and weight can be generated by arbitrary methods. Moreover, the discrete settings under ROI constrains also can be readily supported with few adaptations.

4 EXPERIMENT

To demonstrate the effectiveness of our approach, we present results of both offline experiments and online A/B testing.

4.1 Market Response

4.1.1 Datasets. The proposed semi-black-box response model is tested on two datasets.

- **Breakfast.** It is a public dataset³ and contains sales and promotion records on 4 thousand market-segments (58 products from 77 stores) over 156 weeks. For each market-segment, the sales volume is accumulated weekly, and the market size is approximated by maximum sales in history. Cost on marketing (discount or premium) can be calculated from the difference between *shelf price* and *base price*. Contextual variables include store, product, manufacture, category, sub-category, in-store display, etc.
- **Taopiaopiao.** It is collected from Taopiaopiao⁴, which is one of the largest platforms for selling movie tickets in China. There are 60 thousand market-segments (each day-of-week of 8 thousand cinemas) over 152 days. For each market-segment, the sales volume is accumulated daily, and the market size is provided by government agency⁵. Cost on marketing can be obtained directly, and contextual variables include cinema, day-of-week, city, etc.

³<https://www.dunnhumby.com/careers/engineering/sourcefiles>

⁴<https://www.taopiaopiao.com>

⁵<https://zgdyyp.cn>




Figure 4: The comparison between logit model and our semi-black-box model under different partition ratios.

4.1.2 Setups. As we can see, all contextual variables are indicator variables. In our model, they are firstly represented by one-hot encodings and then concatenated together as input layer of the neural network part. After that, there are 5 fully connected layers with ReLU as the activation function, and the dimension of each layer is 16. In both the baseline (logit model) and the proposed semi-black-box model, all trainable parameters are initialized as zero and updated through gradient decent with Adam rule [18]. The learning rate is 0.01 and all models are trained for 200 epochs.

4.1.3 Results. The results are measured in Relative Mean Absolute Error (RMAE), which is calculated as:

$$\text{RMAE} = \frac{\sum_{i=1}^N |y_i - \hat{y}_i|}{\sum_{i=1}^N \hat{y}_i}, \quad (23)$$

where N is the number of test samples, \hat{y}_i is the true value and y_i is the prediction value. For both Breakfast and Taopiaopiao, the whole datasets are split into training and test sets along business timeline. In order to comprehensively evaluate our method, we measure performances under different partition ratios. In particular, the proportion of dataset selected as training set ranges from 5% to 95%, and all results are plotted in Figure 4. As can be seen, our semi-black-box model significantly outperforms the traditional logit model, especially when the training data is limited. The reason is that our model is able to share information across different segments, and thus data sparseness problem is greatly relieved.

4.2 Budget Allocation

To investigate the behaviors of our budget allocation algorithm, several simulations are conducted.

4.2.1 Setups. We consider a synthetic scenario including $N = 100$ market-segments, where the parameters of market response models are generated randomly. More specifically, the parameters are generated as follow: for $i=1, \dots, N$, $D_i \sim \mathcal{U}(0, 100)$, $a_i \sim \mathcal{U}(-1, 1)$, $b_i \sim$




Figure 5: Convergence processes of three random examples (grouped by column).

$\mathcal{U}(0, 1)$ and $B \sim \mathcal{U}(0, 100N)$, where $\mathcal{U}(l, u)$ represents the uniform distribution on interval (l, u) .

4.2.2 Convergence. The convergence processes of three random examples are shown in Figure 5. We can see that the proposed algorithm converges in a few iterations, usually less than 10 iterations. As $g(\lambda)$ goes to 0, the objective value approaches its maximum value, which confirms the derivation in Section 3.2.

4.2.3 Parameter sensitivity. To study how the estimation bias of a_i and b_i disturbs our algorithm, some artificial biases are simulated on the generated parameters.

Firstly, after adding different level biases to the generated $a_i, b_i (i = 1, \dots, N)$, we get the biased parameters:

$$\hat{a}_i = a_i + \epsilon |a_i|, \quad \hat{b}_i = b_i + \epsilon |b_i|, \quad (24)$$

where ϵ is randomly sampled from $\mathcal{U}(-0.2, 0.2)$ for each group $a_i, b_i (i = 1, \dots, N)$. Secondly, \hat{c}_i is solved under (\hat{a}_i, \hat{b}_i) and (a_i, b_i) respectively. Based on \hat{c}_i , objective value and constraint function (or exceeded budget) are calculated using a_i and b_i :

$$\hat{d}(\hat{c}) = \sum_{i=1}^N D_i q_i(\hat{c}_i; a_i, b_i), \quad \hat{g}(\hat{c}) = \sum_{i=1}^N D_i q_i(\hat{c}_i; a_i, b_i) \hat{c}_i - B. \quad (25)$$

Then we compute the relative error of objective value and exceeded budget as follow:

$$\begin{aligned} \text{Objective error} &= \frac{d(c^*) - \hat{d}(\hat{c})}{d(c^*)}, \\ \text{Exceeded budget} &= \frac{\hat{g}(\hat{c})}{B}. \end{aligned} \quad (26)$$

More than 1000 independent Monte Carlo simulations are conducted, and the results are shown in Figure 6.

Obviously, when the estimation bias increases, both objective error and exceeded budget increase. However, the objective is not very sensitive to parameters, e.g. the objective error is still less than 1% even with 20% estimation bias. On the other hand, although budget is more sensitive to parameters than objective, the exceeded budget is less than 5% when estimation biases of a_i, b_i are less than 10%, 5% respectively. Moreover, the algorithm is less sensitive to a_i than b_i . One reason is that, in logit response models, b_i is the coefficient of decision variable c while a_i is the bias term. Note that estimation bias is limited under a reasonable forecasting accuracy.




Figure 6: Parameter sensitivity. (a) (b) The mean and standard deviation of objective error and exceeded budget with different estimation biases of a . (c) (d) The mean and standard deviation of objective error and exceeded budget with different estimation biases of b .

4.2.4 Discrete settings. The discrete constraints are common in many real-world situations, and we have presented efficient approximation algorithms for solving discrete problems. To evaluate the approximate error of proposed methods, several simulations are conducted. Two approximation algorithms are tested in simulations: 1) Bisection + Dyer-Zemel, which solves the relaxed continuous problem first and then solves the constructed MCKP by Dyer-Zemel; 2) Dyer-Zemel only, which is suggested when the set of options is small. As the distance between two adjacent optional values, which indicates the density in discrete constraints, varies from 0.1 to 8, 100 independent Monte Carlo simulations are run for each distance.

The upper bound of approximate error can be calculated as follow:

$$\text{Approximate Error} = \frac{d_u - d_a}{d_u - d_0}, \quad (27)$$

where d_a is the objective value of approximate solution, d_0 is the objective value with no action (i.e. $c_i = 0$ for every i), and d_u is the maximum objective value of relaxed continuous problem. The detailed simulation results are shown in Table 1, which includes both mean and standard derivation of approximate errors. It is shown that the upper bound of approximate error is very small when the options have a high density (or small discrete distance). For example, when discrete distances are less than 1, the approximate error is no more than 1%. As optional values become sparse, the upper bound of approximate error increases, which is consistent with our intuition. However, the real approximate error, which is unknown due to NP-hardness, is not necessarily significant. Moreover, the budget allocation results of two algorithms are the same in most situations. Therefore, it is reasonable to suggest directly solving problems with discrete constraints by Dyer-Zemel algorithm when the set of options is small and the running time is not the bottleneck.

Table 1: The approximate error upper bound of two algorithms for solving discrete problems.

Distance	Approximate error (mean \pm stddev) (%)	
	Bisection + Dyer-Zemel	Dyer-Zemel only
0.1	0.0297 \pm 0.0337	0.0299 \pm 0.0337
0.5	0.3040 \pm 0.2614	0.3043 \pm 0.2612
1	0.8418 \pm 0.5582	0.8418 \pm 0.5582
2	3.053 \pm 2.247	3.053 \pm 2.247
4	15.14 \pm 10.43	15.09 \pm 10.45
8	49.89 \pm 11.52	49.88 \pm 11.53

4.3 Real-world Scenarios

In order to show the performance of our approach in real-world scenarios, we would like to take the daily marketing in Taopiaopiao as a typical example. Before policy changes, there are several billion RMB spent on marketing each year to compete for market shares (as equal to sales) with other platforms, Maoyan for an example. How to spend money efficiently is a crucial part in the competition.

4.3.1 Setups. In order to avoid price discrimination, the unit cost (price discount or premium) on marketing is determined for each market-segment (i.e. each day-of-week of different cinemas). For instance, the discounts of movie tickets on Sunday for a specific cinema are the same for all Taopiaopiao customers, but the discounts can be different across different cinemas and also can be different from Saturday. The cost upper bound is updated weekly. Based on initial budget, the allocation is firstly calculated at the beginning of the following week. Then the allocation need to be readjusted daily based on the left budget. One reason for budget allocation shifting is the biased forecasting. In addition, the discrete constraints are required here due to customers' preference of 99-ending prices.

We design the following cinema-based randomized A/B testing experiment to form fair comparisons. Firstly, we filter out outliers, including cinemas with extreme high or low market shares in history, or abnormal variance of sales, etc. Secondly, 4 hundred cinemas are randomly picked across different cities as group A, and then a one-to-one mapping approach is conducted to generate group B. In the one-to-one mapping approach, the 1-nearest-neighbor is selected for each cinema in group A based on historical data and contextual features. Thirdly, an A/A testing runs for several days to confirm group A and B have similar sales volumes and marketing responses. Finally, an A/B testing runs for five weeks, where cost upper bounds of group A and B for each week are the same. The budget is uniformly allocated in group A, and allocated based on our approach in group B.

4.3.2 Results. The comparison results on market share are shown in Figure 7, and the comparison results on ROI are shown in Figure 8. First of all, the sales is effectively improved every week and the average improvement is above 6%. Actually, the algorithm is always early terminated by $f(\lambda_l) - f(\lambda_r) \leq \epsilon'$. The reason is that the cost upper bound is given by managers according to their subjective experience, and it is not necessary to spend that much money for negligible improvements. So our approach decides to spend much less money, and thus the improvements on ROI is remarkable (more than 40% averagely). In one word, the proposed




Figure 7: Comparison results on market share.




Figure 8: Comparison results on ROI.

approach can improve sales over 6% with spending 40% less money on marketing by allocating budget into the right market-segments. As this example scenario shows, the data-driven methods can help us to make more informed and effective decisions in online business, where it is difficult for human to monitor and analyze the business environment constantly.

5 RELATED WORK

Marketing is one of the premier components of business management, and it has been studied for decades [22]. In marketing, companies need to allocate whole budget across different market-segments, such as across countries, products, consumption months and so on [12]. There are great real benefits for improving the efficiency of marketing budget allocation, and thus many studies focus on this topic [8, 14, 29]. Recently, online business has become an important part of our daily lives. The environment in online business is very dynamic and complex, which brings great challenges to marketing, especially for companies with large and diversified markets [19, 32]. Many data-driven techniques are developed [1, 10]

to overcome these challenges, including forecasting [3, 37] and optimization [16].

As black-box models, neural networks have recently been successfully applied to many forecasting applications [21, 33, 35]. Through multiple non-linear transformations, they can achieve a high forecasting accuracy. However, there are several gaps between forecasting and decision making [2]. One challenge is that black-box models often fail to identify relationship between the independent variable and objective variable, e.g. how cost on marketing effects sales, due to covariate shifting. A related topic on this problem is causal inference [23, 25]. Another challenge is that optimization on black-box functions is difficult. Therefore, economists prefer using explicit expressions to reveal relationship between the independent variable and objective variable. In the context of sales forecasting, the explicit models are called demand curves. There are many kinds of demand curves, such as linear, log-linear, constant-elasticity, logit and so on [15, 26, 34]. We focus on logit demand curve, which is the most popular response model [26, 28, 36]. However, the capability of fully utilizing data is limited when the logit demand curve is fitted for each market-segment separately. That may cause serious data sparseness problem. One way to relieve data sparseness problem is Thompson sampling [11, 13], where independent variables need to change several time for explorations. In this work, we extend the capability of logit demand curves via neural networks and a semi-black-box model has been proposed to relieve the data sparseness problem.

There are also massive works focusing on solving the marketing (or pricing) problems by formulating them as optimization problems. Ferreira et al. optimize pricing decisions by integer programming, and solve it by linear relaxations [10]. Ito et al. maximize future revenue or profit by binary quadratic programming, and solve it by semi-definite relaxations [16]. Boutilier et al. formulate the budget allocation as a sequential decision problem and solve it by MDP [4]. Staib et al. study the general budget allocation problem from a robust optimization perspective [31]. In this work, the marketing budget allocation is firstly formulated as a non-convex problem. Inspired by [7, 20], whose topics are pricing products with no budget constraints, we reformulate the problem into an equivalent convex optimization problem. Efficient algorithms are designed to solve it in both continuous and discrete settings. It is worth mentioning that a major cause of discrete decision variables is the preference of 99-ending display prices in many situations [30]. The discrete settings can be solved by constructing a Multiple Choice Knapsack Problem (MCKP) from the solution of relaxed continuous problem. Dyer et al. [9] and Zemel et al. [38] independently developed approximation algorithms for MCKP, and the algorithm can be further boosted by pruning [27]. Actually, MCKP has wide applications [17], such as advertising optimization in online marketing [24].

6 CONCLUSION

In this paper, we present a novel unified framework for marketing budget allocation, and there are two sequential steps in the proposed framework. Firstly, the response in each market-segment is forecasted by a semi-black-box model, which extends logit demand curves with the capability of neural networks. The new response model not only can share information across different segments

to relieve data sparseness problem, but also keeps the explicit expression between the independent variable and objective variable. Then budget allocation is formulated as an optimization problem based on learned response models. The optimization problem can be solved efficiently by binary search on a dual variable in a few iterations and the complexity of each iteration is $O(N)$. Thus it is readily to handle large-scale problems. Several business constraints, including cost upper bound, profit lower bound, ROI lower bound, and discrete constraints, are supported in one unified paradigm with minor adaptations. This framework is easy to implement and has been successfully applied to many scenarios in Alibaba Group. The results of both offline experiments and online A/B testing demonstrate its effectiveness. Taking the daily marketing in Taopiaopiao as an example, the proposed approach can improve sales over 6% with spending 40% less money on marketing.

For future works, we are interested in exploring relationship between the market cost and contextual variables in the logit response model. There is also a strong motivation to support boundary constraints on decision variables in the optimization part. Moreover, another interesting topic is to theoretically analyze the approximation error of algorithms for solving the discrete setting.

ACKNOWLEDGMENTS

The authors would like to thank Chao Wang, Tuo Shi, Jun Lang, Yu Dong, Tianxiang Xu, Jun Tan, Zulong Chen, Ling Chen for discussions and supports.

REFERENCES

- [1] Shahriar Akter and Samuel Fosso Wamba. 2016. Big data analytics in E-commerce: a systematic review and agenda for future research. *Electronic Markets* 26, 2 (2016), 173–194.
- [2] Susan Athey. 2017. Beyond prediction: Using big data for policy problems. *Science* 355, 6324 (2017), 483–485.
- [3] Samaneh Beheshti-Kashi, Hamid Reza Karimi, Klaus-Dieter Thoben, Michael Lütjen, and Michael Teucke. 2015. A survey on retail sales forecasting and prediction in fashion markets. *Systems Science & Control Engineering* 3, 1 (2015), 154–161.
- [4] Craig Boutilier and Tyler Lu. 2016. Budget Allocation using Weakly Coupled, Constrained Markov Decision Processes. In *UAI*.
- [5] Stephen Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge university press.
- [6] Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. 1996. On the LambertW function. *Advances in Computational mathematics* 5, 1 (1996), 329–359.
- [7] Lingxiu Dong, Panos Kouvelis, and Zhongjun Tian. 2009. Dynamic pricing and inventory control of substitute products. *Manufacturing & Service Operations Management* 11, 2 (2009), 317–339.
- [8] Peter Doyle and John Saunders. 1990. Multiproduct advertising budgeting. *Marketing Science* 9, 2 (1990), 97–113.
- [9] Martin E Dyer. 1984. An $O(n)$ algorithm for the multiple-choice knapsack linear program. *Mathematical programming* 29, 1 (1984), 57–63.
- [10] Kris Johnson Ferreira, Bin Hong Alex Lee, and David Simchi-Levi. 2015. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* 18, 1 (2015), 69–88.
- [11] Kris Johnson Ferreira, David Simchi-Levi, and He Wang. 2018. Online network revenue management using Thompson sampling. *Operations research* 66, 6 (2018), 1586–1602.
- [12] Marc Fischer, Sönke Albers, Nils Wagner, and Monika Frie. 2009. Dynamic Marketing Budget Allocation across Countries, Products, and Marketing Activities. (2009).
- [13] Ravi Ganti, Matyas Sustik, Quoc Tran, and Brian Seaman. 2018. Thompson Sampling for Dynamic Pricing. *arXiv preprint arXiv:1802.03050* (2018).
- [14] Sunil Gupta and Thomas Steenburgh. 2008. Allocating marketing resources. *Marketing Mix Decisions: New Perspectives and Practices*, Roger A. Kerin and Rob O’ÄZRegan, eds., American Marketing Association, Chicago, IL (2008), 90–105.
- [15] Dominique M Hanssens, Leonard J Parsons, and Randall L Schultz. 2003. *Market response models: Econometric and time series analysis*. Vol. 12. Springer Science & Business Media.
- [16] Shinji Ito and Ryohei Fujimaki. 2017. Optimization beyond prediction: Prescriptive price optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1833–1841.
- [17] Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. The multiple-choice knapsack problem. In *Knapsack Problems*. Springer, 317–347.
- [18] D Kinga and J Ba Adam. 2015. A method for stochastic optimization. In *International Conference on Learning Representations*.
- [19] Kenneth C Laudon, Carol Guercio Traver, et al. 2016. *E-commerce: business, technology, society*.
- [20] Hongmin Li and Woonghee Tim Huh. 2011. Pricing multiple products with the multinomial logit and nested logit models: Concavity and implications. *Manufacturing & Service Operations Management* 13, 4 (2011), 549–563.
- [21] Binbing Liao, Jingqing Zhang, Chao Wu, Douglas McIlwraith, Tong Chen, Shengwen Yang, Yike Guo, and Fei Wu. 2018. Deep Sequence Learning with Auxiliary Information for Traffic Prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 537–546.
- [22] Naresh K Malhotra. 2007. Review of marketing research. In *Review of Marketing Research*. Emerald Group Publishing Limited, v–v.
- [23] Stephen L Morgan and Christopher Winship. 2014. *Counterfactuals and causal inference*. Cambridge University Press.
- [24] Abhishek Pani, S Raghavan, and Mustafa Sahin. 2017. Large-Scale Advertising Portfolio Optimization in Online Marketing. (2017).
- [25] Judea Pearl et al. 2009. Causal inference in statistics: An overview. *Statistics surveys* 3 (2009), 96–146.
- [26] Robert Lewis Phillips. 2005. *Pricing and revenue optimization*. Stanford University Press.
- [27] David Pisinger. 1995. A minimal algorithm for the multiple-choice knapsack problem. *European Journal of Operational Research* 83, 2 (1995), 394–410.
- [28] Huashuai Qu, Ilya O Ryzhov, and Michael C Fu. 2013. Learning logistic demand curves in business-to-business pricing. In *Proceedings of the 2013 Winter Simulation Conference: Simulation: Making Decisions in a Complex World*. IEEE Press, 29–40.
- [29] Werner Reinartz, Jacquelyn S Thomas, and Viswanathan Kumar. 2005. Balancing acquisition and retention resources to maximize customer profitability. *Journal of marketing* 69, 1 (2005), 63–79.
- [30] Robert M Schindler and Thomas M Kibarian. 1996. Increased consumer sales response though use of 99-ending prices. *Journal of Retailing* 72, 2 (1996), 187–199.
- [31] Matthew Staib and Stefanie Jegelka. 2017. Robust Budget Allocation via Continuous Submodular Functions. In *International Conference on Machine Learning*. 3230–3240.
- [32] Judy Strauss and Raymond D Frost. 2016. *E-marketing: Instructor’s Review Copy*. Routledge.
- [33] Souhaib Ben Taieb, Jiafan Yu, Mateus Neves Barreto, and Ram Rajagopal. 2017. Regularization in Hierarchical Time Series Forecasting with Application to Electricity Smart Meter Data. In *AAAI*. 4474–4480.
- [34] Kalyan T Talluri and Garrett J Van Ryzin. 2006. *The theory and practice of revenue management*. Vol. 68. Springer Science & Business Media.
- [35] Yongxin Tong, Yuqiang Chen, Zimu Zhou, Lei Chen, Jie Wang, Qiang Yang, Jieping Ye, and Weifeng Lv. 2017. The simpler the better: a unified approach to predicting original taxi demands based on large-scale online platforms. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1653–1662.
- [36] Ruben van de Geer, Arnoud V den Boer, Christopher Bayliss, Christine Currie, Andria Ellina, Malte Esders, Alwin Haensel, Xiao Lei, Kyle DS Maclean, Antonio Martinez-Sykora, et al. 2018. Dynamic Pricing and Learning with Competition: Insights from the Dynamic Pricing Challenge at the 2017 INFORMS RM & Pricing Conference. *arXiv preprint arXiv:1804.03219* (2018).
- [37] Peng Ye, Julian Qian, Jieying Chen, Chen-hung Wu, Yitong Zhou, Spencer De Mars, Frank Yang, and Li Zhang. 2018. Customized Regression Model for Airbnb Dynamic Pricing. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 932–940.
- [38] Eitan Zemel. 1984. An $O(n)$ Algorithm for the Linear Multiple Choice Knapsack Problem and Related Problems. *Inf. Process. Lett.* 18, 3 (1984), 123–128.

A SUPPLEMENT

A.1 Implementation

In order to improve the reproducibility, we present more details about the implementation.

A.1.1 Software versions. The details of languages, packages and frameworks used in the deployment and experiments are as follow:

- Language: Python 2.7.15, Cython 0.28.5, SQL.

- Packages: NumPy 1.15.4, SciPy 1.1.0, Tensorflow 1.11.0.
- Frameworks: MaxCompute, Spark 2.3.0.

A.1.2 Large scale. When the training set of forecasting is extremely large, batch gradient descent should be used instead of gradient descent, and machines with GPUs are preferred. However, the inference part can be easily paralleled. For the optimization part, our algorithm can be naturally implemented in a distributed framework, such as Spark, where computations are paralleled across market-segments.

A.1.3 Other tips.

- The experimental results on semi-black-box models are quite insensitivity to hyper-parameters in neural networks, and more layers or higher dimensions are not necessary.
- In some extreme cases, it may need to use high-precision computing to calculate q in Eq. (15b) to avoid overflow caused by the exponent. The extreme cases only happen when ϵ is set to be a very small value, but it is not necessary in most situations.
- The pseudocode in page 323 of [17] is a good reference for implementing the Dyer-Zemel algorithm. However, there are two misprints: the condition in step 6 should be $\sum_{i=1}^m w_{ia_i} > c$ and the condition in step 7 should be $\sum_{i=1}^m w_{ia_i} \leq c$.

A.2 Proofs

A.2.1 Theorem 1.

PROOF. The derivative of Eq. (1) is

$$\nabla_c d_i(c) = \frac{D_i b_i \exp\{-(a_i + b_i c)\}}{[1 + \exp\{-(a_i + b_i c)\}]^2}. \quad (28)$$

Let us substitute it into Eq. (3) and we can get:

$$e_i(c) = -\frac{b_i c}{1 + \exp\{a_i + b_i c\}}, \quad (29)$$

then we have $e_i(\hat{c}_i) = e_i(-\frac{a_i}{b_i}) = -\frac{a_i}{2}$. \square

A.2.2 Theorem 2.

PROOF. Note that

$$\frac{dg}{dq_i} = \frac{D_i}{b_i} (1 - a_i + \ln \frac{q_i}{1 - q_i} + \frac{q_i}{1 - q_i}), i = 1, \dots, N, \quad (30)$$

and then

$$\frac{d^2 g}{d^2 q_i} = \frac{D_i}{b_i q_i (1 - q_i)^2}, i = 1, \dots, N. \quad (31)$$

That is $\nabla_{\mathbf{q}}^2 g(\mathbf{q}) = \text{diag}([\frac{D_i}{b_i q_i (1 - q_i)^2}])$. Since $D_i > 0, b_i > 0$ and $0 < q_i < 1$, we have $\nabla_{\mathbf{q}}^2 g(\mathbf{q}) \geq \text{diag}([\frac{27 D_i}{4 b_i}]) > 0$, where equality holds if and only if $q_i = \frac{1}{3}, \forall i$. Thus, $g(\mathbf{q})$ is *strongly convex*. \square

A.2.3 Theorem 3.

PROOF. From Eq. (12d) and $D_i > 0$, we can get:

$$\lambda^* (1 - a_i + \ln \frac{q_i^*}{1 - q_i^*} + \frac{q_i^*}{1 - q_i^*}) = b_i. \quad (32)$$

Since $b_i > 0$, we have $\lambda^* \neq 0$.

Let us define $x_i := \frac{q_i^*}{1 - q_i^*}$ and substitute it into (32), we have

$$\ln x_i + x_i = a_i + \frac{b_i}{\lambda^*} - 1, \quad (33)$$

or equivalently,

$$x_i e^{x_i} = \exp(a_i + \frac{b_i}{\lambda^*} - 1). \quad (34)$$

Since $\exp(a_i + \frac{b_i}{\lambda^*} - 1) > 0$, we have

$$x_i = W(\exp(a_i + \frac{b_i}{\lambda^*} - 1)), \quad (35)$$

where $W(\cdot)$ is the Lambert W function [6]. Then we can obtain

$$q_i^* = \frac{x_i}{x_i + 1} = \frac{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1))}{W(\exp(a_i + \frac{b_i}{\lambda^*} - 1)) + 1}. \quad \square$$

A.2.4 Theorem 4.

PROOF. Firstly, let $x_i := \frac{q_i}{1 - q_i}$. According to

$$\begin{aligned} \frac{d\lambda}{dq_i} &= \frac{d\lambda}{dx_i} \frac{dx_i}{dq_i} \\ &= \frac{-b_i(\frac{1}{x_i} + 1)}{(1 - a_i + \ln x_i + x_i)^2} \frac{1}{(1 - q_i)^2} \\ &= -\frac{\lambda^2}{b_i q_i (1 - q_i)^2}, i = 1, \dots, N, \end{aligned} \quad (36)$$

we have:

$$\frac{dq_i}{d\lambda} = -\frac{b_i q_i (1 - q_i)^2}{\lambda^2}, i = 1, \dots, N. \quad (37)$$

The derivative of $f(\lambda)$ with respect to λ is:

$$\begin{aligned} \frac{df}{d\lambda} &= \sum_{i=1}^N D_i \frac{dq_i}{d\lambda} \\ &= -\sum_{i=1}^N \frac{D_i b_i q_i (1 - q_i)^2}{\lambda^2} < 0, \end{aligned} \quad (38)$$

where $D_i > 0, b_i > 0$ and $0 < q_i < 1$. Thus $f(\lambda)$ is strictly decreasing with respect to λ .

Secondly, based on Eq. (12d) we have:

$$1 - a_i + \ln x_i + x_i = \frac{b_i}{\lambda} \quad (39)$$

Then according to Eq. (30), we have:

$$\frac{dg_i}{dq_i} = \frac{D_i}{\lambda}, i = 1, \dots, N. \quad (40)$$

The derivative of $g(\lambda)$ with respect to λ is:

$$\begin{aligned} \frac{dg}{d\lambda} &= \sum_{i=1}^N \frac{dg_i}{d\lambda} = \sum_{i=1}^N \frac{dg_i}{dq_i} \frac{dq_i}{d\lambda} \\ &= -\sum_{i=1}^N \frac{D_i b_i q_i (1 - q_i)^2}{\lambda^3} < 0, \end{aligned} \quad (41)$$

where $D_i > 0, b_i > 0$ and $0 < q_i < 1$. Thus $g(\lambda)$ is strictly decreasing with respect to λ .

Finally, since $\lim_{\lambda \rightarrow 0} q_i(\lambda) = 1$ and $\lim_{q_i \rightarrow 1} c_i(q_i) = +\infty, i = 1, \dots, N$, we have $\lim_{\lambda \rightarrow 0} g(\lambda) \rightarrow +\infty > 0$. \square

A.2.5 Theorem 5.

PROOF. Firstly, $\nabla_{\mathbf{q}}^2 g'(\mathbf{q}) = \text{diag}([\frac{RD_i}{b_i q_i (1-q_i)^2}])$. Since $R > 0, D_i > 0, b_i > 0$ and $0 < q_i < 1$, we have $\nabla_{\mathbf{q}}^2 g'(\mathbf{q}) \geq \text{diag}([\frac{27RD_i}{4b_i}]) > 0$, where equality holds if and only if $q_i = \frac{1}{3}, \forall i$. Thus, $g'(\mathbf{q})$ is *strongly convex*.

Secondly, the derivative of $g'(\lambda)$ with respect to λ is:

$$\frac{dg'}{d\lambda} = - \sum_{i=1}^N \frac{D_i b_i q_i (1-q_i)^2}{R \lambda^3} < 0, \quad (42)$$

where $R > 0, D_i > 0, b_i > 0$ and $0 < q_i < 1$. Thus $g'(\lambda)$ is strictly decreasing with respect to λ .

Finally, since $\lim_{\lambda \rightarrow 0} q_i(\lambda) = 1$ and $\lim_{q_i \rightarrow 1} c_i(q_i) = +\infty, i = 1, \dots, N$, we have $\lim_{\lambda \rightarrow 0} g'(\lambda) \rightarrow +\infty > 0$. \square