Traffic Shaping in E-Commercial Search Engine: Multi-Objective Online Welfare Maximization

Liucheng Sun[1,2], Chenwei Weng[1], Chengfu Huo[1], Weijun Ren[1], Guochuan Zhang[2], Xin Li[1]
[1]Alibaba Group , [2]Zhejiang University
liucheng.slc@alibaba-inc.com, wengchenwei.pt@alibaba-inc.com, chengfu.huocf@alibaba-inc.com
afei@alibaba-inc.com, zgc@zju.edu.cn, xin.l@alibaba-inc.com

Abstract

The e-commercial search engine is the primary gateway for customers to find desired products and engage in online shopping. Besides displaying items to optimize for a single objective (i.e., relevance), ranking items needs to satisfy some other business requirements in practice. Recently, traffic shaping was introduced to incorporate multiple objectives in a constrained optimization framework. However, many practical business requirements can not explicitly represented by linear constraints as in the existing work, and this may limit the scalability of their framework. This paper presents a unified framework from the aspect of multi-objective welfare maximization where we regard all business requirements as objectives to optimize. Our framework can naturally incorporate a wide range of application-driven requirements. In addition to formulating the problem, we design an online traffic splitting algorithm that allows us to flexibly adjust the priorities of different objectives, and it has rigorous theoretical guarantees over the adversarial scenario. We also run experiments on both synthetic and real-world datasets to validate our algorithms.

Introduction

With the rapid growth of e-commerce, online product search has emerged as a popular and effective paradigm for users to find desired products and engage in online shopping. When a user enters a query, search engine returns a page consists of relevant items. To maximize user engagement, it is practical and common to predict the relevance scores by a well-designed machine learning model (or multiple models) and select the top K items with the highest relevance to the query as the final ranking result.

Proxies such as click-through rate (CTR) or conversion rate (CVR) are used for computing relevance scores. While such a single metric-focused ranking method drives up short-term engagement, it ignores long-term sustainability. A click or purchase is only a part of a user's journey in the platform and subsequent downstream utilities such as retention, re-purchase and browsing time-spent are also important. On the other hand, the e-commerce platform is a two-sided marketplace that have customers not only on the demand side (e.g. users), but also on the supply side (e.g. sellers), but simply maximizing users' immediate engagement may fail to satisfy many sellers' demands. For example, since the items from top sellers are usually temporarily more attractable (higher predicted CTR and CVR) than that of new sellers, new sellers may be hard to survive. The attention given to sellers as a result of predicted relevance is often vastly unequal. A relatively small group of superstar sellers receive a large portion of attention in the platform while the majority of sellers in the long tail receives very little. In order to boost market vitality, the e-commerce platform, as a distribution channel to match items/sellers and users, needs to balance different stakeholders' short/long-term desires.

Therefore, the traffic shaping is crucial and has received significant attention by scholars and practitioners, which aims at re-ranking the candidate items to consider additional business requirements. Shaping traffic in such scenarios often entails balancing multiple conflicting objectives and becomes a non-trivial task. Another challenge is the online aspect of the problem: each time a query arrives dynamically, one needs to make an instant and irrevocable decision on what items to display while satisfying global business requirements over the entire set of queries. In this paper, we aim to design a smoothly online matching policy (between items and queries) that balances the trade-off among multiple objectives.

Before we proceed, we very briefly review the existing literature on the traffic shaping problem and position our contributions. Most previous studies consider the problem of maximizing myopic objectives like CTR but subject to some constraints that are intuitive and help in driving long-term value the platform seeks to attain (Agarwal et al. 2011, 2012; Chen and Wang 2015; Gupta et al. 2016; Ding, Govindaraj, and Vishwanathan 2019). In this way, the traffic shaping is regarded as an online constrained optimization problem. To design an online algorithm, they draw on ideas from the primal-dual method by incorporating a dual-learning phase followed by an online decision-making phase (Devanur and Hayes 2009; Agrawal, Wang, and Ye 2014; Chen and Wang 2015). The algorithm is empirically shown to be effective under the assumption that the traffic distribution is stationary. However, the stationariness is generally not true due to the traffic spike as pointed out by (Esfandiari, Korula, and Mirrokni 2018), and the algorithm may suffer severely degenerated performance. Besides, most of them do not present theoretical guarantees on the solution quality except (Shah, Soni, and Chevalier 2017). Moreover, many practical business requirements can not explicitly represented by linear constraints as in all the aforementioned prior work, and this may limit the scalability of their framework.

In this paper, we propose a unified framework from the aspect of multi-objective optimization for dealing with traffic shaping problem. The key insight is that instead of imposing the business constraints, we regard all business requirements as objective to optimize. The designed algorithm allows us to flexibly adjust the priorities of different objectives. It is robust with any traffic distributions and has rigorous theoretical guarantees. Our framework can naturally incorporate a wide range of application-driven requirements. At the time of writing, the method was fully deployed on a large e-commerce website for the traffic shaping task. To summarize, this paper provides three-fold contributions.

- First, we propose a unified framework for dealing with the traffic shaping problem in e-commercial search engine, where users arrive in an online manner and we wish to make decisions that optimize for multiple objectives. In particular, we aim at maximizing the sum of relevance scores, the delivery of guaranteed impressions/clicks and the number of sold items.
- Second, we develop a provably efficient traffic splitting algorithm and provide formal mathematical guarantees. To do so, we show that the objective functions are submodular (linear, budget-additive and probabilistic coverage, respectively) and exploit the primal-dual techniques to design and analyze the online algorithm. In particular, under some mild assumptions, we prove the following main theorem.

  Theorem 1. *There is a $(p_1, p_2(1-e^{-\frac{1}{p_2}}), p_3/(1+p_3))$-competitive algorithm for our tri-objective traffic shaping problem, for any given $p_1, p_2, p_3 \geq 0$ with $p_1+p_2+p_3 \leq 1$.*
- Third, extensive experiments are conducted. The simulation results on a synthetic dataset validate the effectiveness of our proposed algorithm and the A/B testing results in the real-world production environment exhibit the effectiveness and practicability of the method.

The remainder of the paper is structured as follows. Section 2 gives some preliminaries. Section 3 introduces the problem statement and formulation. The basic algorithm and the improved algorithm are presented in Section 4. Section 5 reports our experiments followed by additional related work in Section 6, and Section 7 concludes the paper.

Preliminaries

In this section, we present some definitions and useful tools that are necessary for our explanation and analysis throughout the paper.

Submodular function. If Ω is a finite set, a *submodular* function is a valuation function $w : 2^{|\Omega|} \to \mathbb{R}$, where $2^{|\Omega|}$ denotes the power set of Ω, which satisfied the following condition:

$$w(S \cup e) - w(S) \geq w(T \cup e) - w(T) \quad (1)$$

for every $S, T \subseteq \Omega$ with $S \subseteq T$ and every $e \in \Omega \backslash T$. Moreover, we say w is *monotone* if for every $S \subseteq T \subseteq \Omega$, $w(S) \leq w(T)$.

There are other equivalent definitions but the above formula intuitively shows that submodular functions capture a natural diminishing returns property, which makes them suitable for many applications. Therefore it is natural to study submodular optimization, and there is a large amount of work on minimizing/maximizing submodular functions subject to some constraints. Among them, we are concerned with the *submodular welfare maximization* in this paper.

Welfare maximization and its variants. In the welfare maximization problem (sometimes also referred to as "combinatorial auctions"), the goal is to allocate $|\Omega|$ items to n agents with valuation functions $w_i : 2^{|\Omega|} \to \mathbb{R}^+$ in a way that maximizes $\sum_{i=1}^{n} w_i(S_i)$,where S_i is the set of items allocated to agent i (satisfying $S_i \cap S_j = \emptyset$ for $i \neq j$ in most cases).

In the online version of the problem, items arrive one by one and we have to allocate each item when it arrives, knowing only the agents' valuations on the items that have arrived so far. If the valuation functions are monotone and submodular, we call it the *online submodular welfare maximization* (online SWM) problem. Moreover, if each agent is associated multiple valuation functions that we hope to optimize simultaneously, it becomes *multi-objective online submodular welfare maximization* problem. In later sections, we will see that our problem falls into this variant. Besides, in our problem, an item is allowed to allocate to multiple agents, and we call this feature as "multiple allocation".

Competitive ratio. The competitive ratio is a commonly-used metric to evaluate the performance of online algorithms. Consider an online maximization problem for example. Let $ALG(I) = \mathbb{E}[ALG(I)]$ denote the expected performance of ALG on a concrete instance I, where the expectation is taken over the arrival sequence I. Let $OPT(I) = \mathbb{E}[OPT(I)]$ denote the expected offline optimal, where $OPT(I)$ refers to the optimal value after we observe the full arrival sequence I. Then, the competitive ratio is defined as $\min_I \frac{ALG(I)}{OPT(I)}$. In our paper, we conduct competitive ratio analysis on triple objectives.

Primal-dual framework. In the online optimization literature, it is common to use the primal-dual framework to design or analyze the algorithm. The idea is briefly stated as follows. We use a Linear Program (LP) to upper bound OPT and note its corresponding dual problem. Let P and D be the values of the objective functions of the primal and dual solutions produced, respectively. Initially, $P = D = 0$. We focus on a single iteration of the algorithm and denote by ΔP and ΔD the change in the primal and dual cost, respectively. When the algorithm runs, we maintain a pair of primal and dual solutions to keep the primal and dual feasible and $\Delta P \geq \Gamma \cdot \Delta D$ all the time. As a result, we have $ALG \geq \Gamma \cdot D \geq \Gamma \cdot OPT$ where the second inequality follows the weak duality property. Hence the competitive ratio is $\frac{ALG}{OPT} \geq \Gamma$. For more details, we refer the interested readers to (Buchbinder and Naor 2007; Devanur, Jain, and Kleinberg 2013).

Problem Statement and Formulation

As established in Section 1, when a query arrives, we need to return a list of items to optimize several objectives. In this section, we formally describe the details of the traffic shaping problem that arises in the e-commercial search engine and present the offline tri-objective optimization formulation.

Notations. Let $\mathcal{I}$, $\mathcal{J}$ denote the collections of items and queries respectively. For any query $j \in \mathcal{J}$, let $N(j)$ be the set of items that are retrieved for ranking by the search engine. Similarly, for any item $i \in \mathcal{I}$, let $N(i)$ denote the set of matched queries. We use S_i to represent the query set that allocated to item i. Other notations will be introduced when needed.

In the following, we elaborate on three typically objectives considered in the e-commerce platform.

- **Relevance.** To optimize the short-term user engagement, the first and essential objective is to maximize the sum of relevance scores. Let r_{ij} be the relevance score between item i and query j predicted by the machine learning model [1]. The objective can be written as:

$$\max_{S} \sum_{i \in \mathcal{I}} f_i(S_i) = \sum_{i \in \mathcal{I}} \sum_{j \in S_i} r_{ij}. \tag{2}$$

- **Guaranteed impressions/clicks.** As we state before, in practice there are often other business requirements in place. The most common one adopted in industry is the the minimum number of impressions certain items are served with, or in some cases the minimum number of clicks obtained by certain items. This requirement is intuitive and easy to specify from a business perspective: it may protect the interests of premium or new sellers in e-commerce. To accommodate the requirement in our unified multi-objective optimization framework, we claim that it can be expressed as an objective instead of a constraint in the literature. Since the study on guaranteed impressions is a special case of that on guaranteed clicks (by supposing that every exposed item will be clicked), we focus on the latter. We use C_i to denote the target number of clicks of item i and $\mathcal{I}_c$ the set of guaranteed items. Let v_{ij} be the click probability. The second objective can be formulated as:

$$\max_{S} \sum_{i \in \mathcal{I}_c} g_i(S_i) = \sum_{i \in \mathcal{I}_c} \min\{\sum_{j \in S_i} v_{ij}, C_i\}. \tag{3}$$

  This means that we can not benefit from the extra clicks.

- **Number of sold items.** We present another crucial business requirement in e-commerce: maximizing the number of sold items. There are usually a huge number of items on the platform, but only a few of them have sales. One key metric to evaluate the prosperity of the market is the total number of sold items. Once the new items are sold smoothly in the platform, sellers are more active to publish new products and users are more likely to browse. Thus it creates a virtuous circle and is beneficial to the sustainability of the platform. We use u_{ij} to denote the probability that query (user) j will buy item i and $\mathcal{I}_s$ the set of items with no sales that we hope to be purchased. Note that in practice not all items with no sales are included and the target items are selected due to some business considerations. Assuming that the queries are independent with each other, the probability that item $i \in \mathcal{I}_s$ is sold at least once over a set of queries S_i is:

$$h_i(S_i) = Pr_i(sold|S_i) = 1 - \prod_{j \in S_i}(1 - u_{ij}). \tag{4}$$

  Therefore the third objective can be written as

$$\max_{S} \sum_{i \in \mathcal{I}_s} h_i(S_i) = \sum_{i \in \mathcal{I}_s}(1 - \prod_{j \in S_i}(1 - u_{ij})). \tag{5}$$

We further assume that when query j arrives, a list of b_j distinct items is asked to display in the slots. Let $c_j(S_i)$ denote whether item j is included in set S_i. Put all the pieces together, we get the following ultimate multi-objective optimization formulation:

$$\begin{aligned} \max_{S} \quad & \sum_{i \in \mathcal{I}} f_i(S_i), \sum_{i \in \mathcal{I}} g_i(S_i), \sum_{i \in \mathcal{I}} h_i(S_i) \\ s.t. \quad & \sum_{j \in N(i)} c_j(S_i) = b_j, \qquad \forall j \in \mathcal{J}. \end{aligned} \tag{6}$$

Note that for item $i \in \mathcal{I} \backslash \mathcal{I}_c$, $g_i(S_i) = 0$ and for item $i \in \mathcal{I} \backslash \mathcal{I}_s$, $h_i(S_i) = 0$.

Multi-objective SWM problem with multiple allocation. Finally, we show that the above problem happens to be a multi-objective SWM problem with multiple allocation. The queries and users here are "agents" and "items" respectively defined in the previous section. Each item $i \in \mathcal{I}$ is associated with three functions which evaluate the values of set of queries exposed to it. We verify that all three valuation functions are submodular:

- $f_i(S_i) = \sum_{j \in S_i} r_{ij}$. It is a linear function, known as being submodular(modular).
- $g_i(S_i) = \min\{\sum_{j \in S_i} v_{ij}, C_i\}$. It is a budget-additive function, a typical submodular function.
- $h_i(S_i) = 1 - \prod_{j \in S_i}(1 - u_{ij})$. We call it the *probabilistic coverage function*. It is submodular by definition since it satisfies the condition (without loss of generality, the subscript i is omitted):

$$\begin{aligned} & h(S' \cup j') - h(S') \\ = & \prod_{j \in S'}(1 - u_j) - \prod_{j \in S'}(1 - u_j)(1 - u_{j'}) \\ = & \prod_{j \in S'}(1 - u_j)u_{j'} \\ \leq & \prod_{j \in S}(1 - u_j)u_{j'} \\ = & h(S \cup j') - h(S) \end{aligned} \tag{7}$$

[1] In this paper, we make the simplifying assumption that the predictions are "perfect".

where $S \subseteq S'$ and $j' \notin S' \backslash S$.

Now the only distinction between our problem and the standard multi-objective SWM problem is that a query is allowed to allocate to multiple items in different slots since $b_j > 1$ in most real-world cases. We will show in the subsequent section that the idea for dealing with the single slot setting remains in force by some slight modifications since the position bias and item mutual influence are beyond the scope of this paper.

Traffic Splitting Algorithm

Greedy Traffic Splitting Algorithm (GTSA)

To motivate our idea, we begin with the general multi-objective online SWM with multiple allocation (without considering the particularities of objective functions). The idea of the algorithm is simple but effective: we split the incoming traffic to optimize for different objectives respectively, and the optimization method here is the greedy algorithm which is widely adopted in the submodular optimization. The approach is called the Greedy Traffic Splitting Algorithm (GTSA) and its details are shown in Algorithm 1.

Algorithm 1: Greedy Traffic Splitting Algorithm

Input: $p_1, p_2, p_3 > 0$ with $p1 + p2 + p3 = 1$.
Output: Allocation plan.
Initialize $S_i = \emptyset$ for all $i \in \mathcal{I}$.
while *query j arrives* **do**
  $T = \emptyset$.
  for $k = 1$ *to* b_j **do**
    With probability p_1,
    $i^* = \arg\max_{i \in N(j) \backslash T}[f_i(S_i \cup j) - f_i(S_i)]$.
    With probability p_2,
    $i^* = \arg\max_{i \in N(j) \backslash T}[g_i(S_i \cup j) - g_i(S_i)]$.
    With probability p_3,
    $i^* = \arg\max_{i \in N(j) \backslash T}[h_i(S_i \cup j) - h_i(S_i)]$.
    $S_{i^*} = S_{i^*} \cup \{j\}$.
    $T = T \cup \{i^*\}$.
  end
  Display the items from set T.
end

To obtain the theoretical guarantees of the algorithm, we consider the simplified *Probabilistic Greedy Algorithm* (PGA) as described in Lemma 1.

Lemma 1. *When a query arrives, pass it to the greedy algorithm for allocation in each slot with probability p, and leave it unmatched otherwise. This Probabilistic Greedy Algorithm (PGA) is $\frac{p}{1+p}$-competitive for the online SWM with multiple allocation.*

Proof. Let A_i denote the query set that allocated to item i by the algorithm and thus $w_i(A_i)$ is the actual cumulative gain of item i at the end of the algorithm. Similarly, let O_i be the query set that allocated to item i in the optimal solution. We use $gain_j$ to represent the marginal gain by the algorithm from allocating query j upon its arrival and $\overline{gain}_j = \frac{gain_j}{b_j}$ the average gain. Note that $\overline{gain}_j$ is a random variable, so the expected average marginal gain of allocating query i is $p\mathbb{E}[\overline{gain}_j]$. We first prove that the following key inequality holds for any $i \in \mathcal{I}$ and $j \in \mathcal{J}$:

$$\overline{gain}_j \geq w_i(A_i \cup j) - w_i(A_i). \tag{8}$$

Let A_i^t denote the query set allocated to item i at iteration t (before query j arrives), thus by definition we obtain:

$$\overline{gain}_j = \frac{1}{b_j} \sum_{k=1}^{b_j} w_{i_k}(A_{i_k}^t \cup j) - w_{i_k}(A_{i_k}^t). \tag{9}$$

For the first case, suppose that $i \notin \{i_1, i_2, ..., i_{b_j}\}$ (i.e. item i is not displayed to query j), then for any $1 \leq k \leq b_j$, we have:

$$\begin{aligned} w_{i_k}(A_{i_k}^t \cup j) - w_{i_k}(A_{i_k}^t) &\geq w_i(A_i^t \cup j) - w_i(A_i^t) \\ &\geq w_i(A_i \cup j) - w_i(A_i) \end{aligned} \tag{10}$$

The first inequality follows since our algorithm is greedy and the second inequality follows from the submodularity. For the second case, suppose that $i \in \{i_1, i_2, ..., i_{b_j}\}$, in other words, item i and query j is matched by the algorithm ($j \in A_i$). This implies:

$$w_i(A_i \cup j) - w_i(A_i) = 0. \tag{11}$$

Combining all of the above inequalities and equations, the key inequality (8) holds. By taking the expected value of both sides of inequality (8), we have:

$$p\mathbb{E}[\overline{gain}_j] \geq \mathbb{E}[w_i(A_i \cup j) - w_i(A_i)]. \tag{12}$$

Furthermore, we can establish the relationship between OPT and ALG by linearity of expectation:

$$\begin{aligned} \mathbb{E}[ALG] &= \sum_{i \in \mathcal{I}} \sum_{j \in O_i} p\mathbb{E}[\overline{gain}_j] \\ &\geq \sum_{i \in \mathcal{I}} \sum_{j \in O_i} p\mathbb{E}[w_i(A_i \cup j) - w_i(A_i)] \\ &\geq \sum_{i \in \mathcal{I}} p\mathbb{E}[w_i(A_i \cup O_i) - w_i(A_i)] \\ &\geq \sum_{i \in \mathcal{I}} p\mathbb{E}[w_i(O_i) - w_i(A_i)] \\ &= p\mathbb{E}[OPT] - p\mathbb{E}[ALG] \end{aligned} \tag{13}$$

where the second inequality follows from the submodularity, and the last inequality from the monotonicity. Therefore we obtain the competitive ratio:

$$\frac{\mathbb{E}[ALG]}{\mathbb{E}[OPT]} \geq \frac{p}{1+p}. \tag{14}$$

□

The main theorem of this subsection follows immediately.

Theorem 2. *The GTSA is $(\frac{p_1}{1+p_1}, \frac{p_2}{1+p_2}, \frac{p_3}{1+p_3})$-competitive for the general tri-objective online SWM with multiple allocation, for any given p_1, p_2, p_3 with $p_1 + p_2 + p_3 \leq 1$.*

Improved Traffic Splitting Algorithm (ITSA)

The GTSA can be further improved from both analysis and design, because the first and second objective functions in our problem are a little special. In the following, we first prove that the PGA is indeed better than $\frac{1}{1+p}$-competitive for the first objective (linear valuation functions) and then present a tailored-made algorithm for the second objective (budget-additive valuation functions).

Probabilistic Greedy Algorithm for the First Objective

Lemma 2. *When a query arrives, pass it to the greedy algorithm for allocation in each slot with probability* p*, and leave it unmatched otherwise. The Probabilistic Greedy Algorithm is* p*-competitive for the online SWM with multiple allocation and linear valuation functions.*

Proof. It is easy to verify that greedy is optimal for online SWM with linear valuation functions (even in the multiple allocation scenario), which implies:

$$\begin{aligned}\mathbb{E}[ALG] &= \sum_{i\in\mathcal{I}}\sum_{j\in O_i} p\mathbb{E}[\overline{gain}_j] \\ &= \sum_{i\in\mathcal{I}} p\mathbb{E}[f_i(O_i)] = p\mathbb{E}[OPT].\end{aligned} \tag{15}$$

Therefore, we obtain:

$$\frac{\mathbb{E}[ALG]}{\mathbb{E}[OPT]} \geq p. \tag{16}$$

□

Probabilistic Balance Algorithm for The Second Objective

To design a more smart algorithm for the second objective, we leverage the primal-dual framework which is widely used for the online optimization (as introduced in Section 2). We note the following configuration LP:

$$\begin{aligned}\max_x \quad & \sum_{i,S} x_{iS} g_i(S) \\ s.t. \quad & \sum_S x_{iS} \leq 1, \qquad \forall i \in \mathcal{I} \\ & \sum_{i\in N(j)}\sum_{S\ni j} p x_{iS} \leq b_j, \qquad \forall j \in \mathcal{J} \\ & x_{iS} \geq 0, \qquad \forall i \in I, j \in J\end{aligned} \tag{17}$$

Lemma 3. *The optimal value to LP (17) is a valid upper bound of the expected gain achieved by the offline optimal solution.*

The proof of the above lemma is already quite standard (Bansal et al. 2012; Zhao et al. 2019). Note that it suffices to show that any allocation strategy satisfies the constraints of the LP (17) in expectation. Let α_i and β_j be the associated dual variables, then the corresponding dual problem is:

$$\begin{aligned}\min_{\alpha,\beta} \quad & \sum_{i\in\mathcal{I}} \alpha_i + \sum_{j\in\mathcal{J}} b_j\beta_j \\ s.t. \quad & \alpha_i + p\sum_{j\in S}\beta_j \geq g_i(S), \qquad \forall i \in \mathcal{I}, S \\ & \alpha \geq 0, \beta_j \geq 0, \qquad \forall i \in \mathcal{I}, j \in \mathcal{J}.\end{aligned} \tag{18}$$

Following the primal-dual framework, we design the *Probabilistic Balance Algorithm* (PBA) as shown in Algorithm 2.

Algorithm 2: Probabilistic Balance Algorithm

Fix a non-decreasing function $\phi : [0,1] \to [0,1]$, $0 < p \leq 1$ and $\Gamma > 0$.
Let $y_i = \frac{\sum_{j\in S_i} v_{ij}}{C_i}$ denote the fraction of item i's capacity consumed.
Initialize $\alpha_i = 0, \beta_j = 0$ and $y_i = 0$ for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$.
while *query j arrives* **do**
  $T = \emptyset$.
  $T_{virtual} = \emptyset$.
  for $k = 1$ *to* b_j **do**
    $i^* = \arg\max_{i\in N(j)\setminus T_{virtual}} v_{ij}(1-\phi(y_i))$.
    $T_{virtual} = T_{virtual} \cup \{i^*\}$.
    α_{i^*} is incremented by $v_{ij}\phi(y_{i^*})/\Gamma$.
    With probability p:
    1) $T = T \cup \{i^*\}$;
    2) $y_{i^*} = y_{i^*} + v_{i^*j}/C_{i^*}$.
  end
  $\beta_j = \max_{i\in T} v_{ij}(1-\phi(y_i))/\Gamma$.
  Return item set T for display.
end

By choosing ϕ and Γ carefully, we get the following lemma that provides the theoretical guarantee on the competitive ratio.

Lemma 4. *Set* $\phi(y) = e^{\frac{1}{p}(y-1)}$ *and* $\Gamma = p(1-e^{-\frac{1}{p}})$*, then the Probabilistic Balance Algorithm is* Γ*-competitive under the assumption that* $\max_{i,j}\frac{v_{ij}}{C_i} \to 0$.

Proof. According to the principle of the primal-dual framework described in Section 2, we need to prove that $\Delta P \geq \Gamma\Delta D$ and the dual is feasible. The former is easy to verify: at each iteration (when query j arrives), we have $\Delta P = \sum_{i\in T_{virtual}} v_{ij}$ and

$$\begin{aligned}\Gamma\Delta D = \Gamma\sum_{i\in T_{virtual}}\Delta\alpha_i + b_j\beta_j &= \Gamma\sum_{i\in T_{virtual}}(\Delta\alpha_i + \beta_j) \\ &\leq \sum_{i\in T_{virtual}} v_{ij}(\phi(y_i) + 1 - \phi(y_i)) = \Delta P.\end{aligned} \tag{19}$$

Then we only need to prove the feasibility of the dual. Since $\frac{v_{ij}}{C_i} \to 0$ by assumption, α_i can be approximated as:

$$\alpha_i = C_i\int_0^{y_i}\phi(y)dy/\Gamma. \tag{20}$$

We have $\beta_j = \max_{i\in T_{virtual}} v_{ij}(1-\phi(y_i))/\Gamma$, thus

$$\begin{aligned}\alpha_i + p\sum_{j\in S}\beta_j &\geq C_i\int_0^{y_i}\phi(y)dy/\Gamma + p\sum_{j\in S} v_{ij}(1-\phi(y_i))/\Gamma \\ &\geq \min\{\sum_{j\in S} v_{ij}, C_i\}(\int_0^{y_i}\phi(y)dy + p - p\phi(y_i))/\Gamma.\end{aligned} \tag{21}$$

Note that by setting $\phi(y) = e^{\frac{1}{p}(y-1)}$ and $\Gamma = p(1 - e^{-\frac{1}{p}})$, we have the identical equation $(\int_0^{y_i} \phi(y)dy + p - p\phi(y_i))/\Gamma = 1$ for any $0 \leq y_i \leq 1$. This indicates the feasibility of the dual problem and the algorithm is Γ-competitive. □

Therefore, we can obtain the Improved Traffic Splitting Algorithm (ITSA) as shown in Algorithm 3 and Theorem 1 presented in Section 1 holds immediately. We also depict the competitive ratios on our three different objectives with respect to the value of p.

Algorithm 3: Improved Traffic Splitting Algorithm

Input: $p_1, p_2, p_3 > 0$ with $p1 + p2 + p3 = 1$.
Output: Allocation plan.
Initialize $S_i = \emptyset$ for all $i \in \mathcal{I}$.
while *query* j *arrives* **do**
  $T = \emptyset$.
  for $k = 1$ *to* b_j **do**
    With probability p_1,
    $i^* = \arg\max_{i \in N(j) \setminus T}[f_i(S_i \cup j) - f_i(S_i)]$.
    With probability p_2,
    pass the query to the PBA where $b_j = 1$ and get $i^* \in T_{virtual}$.
    With probability p_3,
    $i^* = \arg\max_{i \in N(j) \setminus T}[h_i(S_i \cup j) - h_i(S_i)]$.
    $S_{i^*} = S_{i^*} \cup \{j\}$.
    $T = T \cup \{i^*\}$.
  end
  Display the items from set T.
end

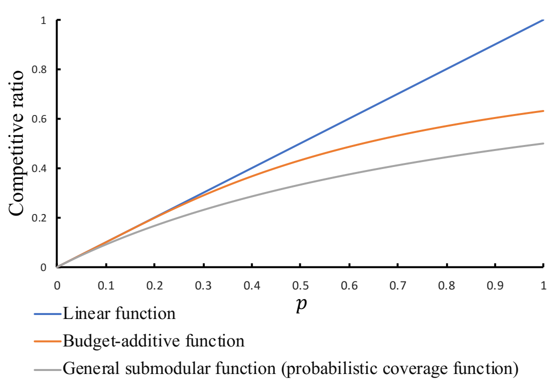

Figure 1: Competitive ratios on different valuation functions with respect to the value of p. The blue, orange and grey lines represent the competitive ratio of the linear function, budget-additive function and general submodular function (probabilistic coverage function) respectively.

Experiments

To demonstrate the feasibility and effectiveness of our proposed method, we conduct a suite of experiments including offline simulation and online production A/B test.

Offline Simulation

Synthetic Dataset To generate a synthetic dataset, we fix $|\mathcal{I}| = 10,000$ and $|\mathcal{J}| = 5,000$. The items are divided into two parts: 2,000 mature items and 8,000 new items to approximately simulate the real environment. For each query j, 200 randomly sampled items are retrieved for ranking, and b_j is sampled from a uniform distribution $U(3, 50)$. For each item-query pair (i, j), the relevance score r_{ij} is drawn from the beta distributions $Beta(3, 2)$ and $Beta(2, 3)$; the click probabilities are drawn from the uniform distributions $U(0.1, 0.3)$ and $U(0, 0.2)$; the conversion probabilities follow the uniform distributions $U(0, 0.01)$ and $U(0, 0.005)$ for mature and new items respectively. Then the purchase probability can be computed by multiplication of the click probability and conversion probability. We set $|\mathcal{I}_c| = 1,000$ and $|\mathcal{I}_s| = 1,000$. The new items in $\mathcal{I}_c$ are randomly sampled, and $|\mathcal{I}_s|$ is composed of 500 mature items and 500 new items (different from items in $\mathcal{I}_c$). The target numbers of clicks are 18 and 2 for mature and new items respectively.

Algorithm We compare three algorithms in the experiment. The first is the relevance-focused method in which we only optimize for a single objective: the sum of the relevance scores. The second is the GTSA, where we fix $p_1 = 0.9$ since the relevance is always the fundamental goal in the search engine. To verify the influence of different values of p, we change the value of p_2 from 0 to 0.1 with the step size 0.01 and $p_3 = 1 - p_1 - p_2$. The third one is the ITSA, the parameter settings are the same as that in the GTSA.

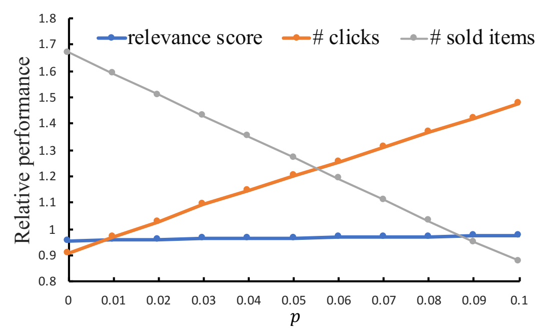

Figure 2: Comparison of the baseline and GTSA on three objectives. The relative performance is computed through the division of the objective function values of the GTSA over the baseline. The blue, orange and grey lines represent the relative performance of the sum of relevance scores, number of clicks on items in $\mathcal{I}_c$ and number of sold items in $\mathcal{I}_s$ respectively.

Results We take the relevance-focused method as the baseline, and focus on the relevance performance of the GTSA and ITSA compared to it. The result of the GTSA is shown in Figure 2. We can see that the performance of the relevance is always above 0.95 with only less than 5% drop in the relevance score. It is better than the strict competitive ratio 0.9 since the exposure of any item will contribute to the total relevance. As expected, the number of clicks of items in $\mathcal{I}_c$ increase with p_2 while the number of sold new items

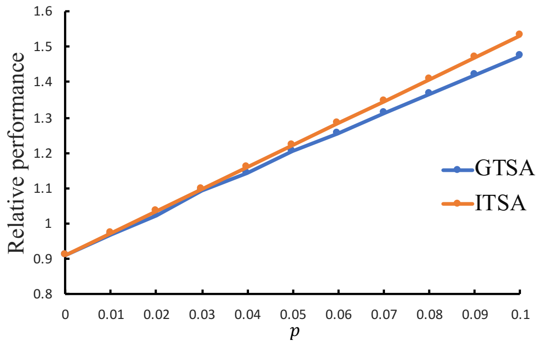

Figure 3: Comparison of the GTSA and ITSA on the second objective. The relative performance is computed through the division of the objective function values of the GTSA/ITSA over the baseline. The blue and orange lines represent the relative performance of the GTSA and ITSA respectively.

	50% Treatment Experimental items	50% Control Experimental items
50% Treatment Queries	B11	B12
50% Control Queries	B21	B22

Table 1: An illustration of the A/B testing design.

in $\mathcal{I}_s$ decreases with p_2. It is interesting to see that when p_2 or p_3 is near to zero, the corresponding performances are worse than the baseline. This is because the second and third objectives are totally conflicting (i.e. they are optimizing for disjoint items).

In addition, we compare the performance of the GTSA and ITSA in terms of the second objective since the other two are treated in the same way. The result is shown in Figure 3. We see that when p_2 is small, they have the similar performance. As p_2 becomes larger, the ITSA outperforms GTSA gradually. This is consistent with the comparison result of competitive ratios shown in Figure 1.

Online A/B Test

The method in this paper (ITSA) has been deployed on a large e-commerce search engine where about 20 million queries and 3 million exposed items involved every day. Besides maximizing the relevance, the platform hopes that more new items (20,000 items are selected) can be sold and some other specified items (20,000 items are selected) can get a certain number of impressions. We split the queries and experimental items equally into two groups respectively and the design of the A/B test is illustrated in table 1. "Treatment" means that we run our algorithm while "Control" means that we follow the traditional relevance-focused method. We set $p_1 = 0.96, p_2 = 0.02, p_3 = 0.02$ and the experiment lasts for a week.

The comparison method is as follows:

- Impact on the relevance: comparing B11+B12 with B21 + B22.
- Impact on the experimental items: comparing B11 + B21 with B12 + B22.

The results are reported as follows. In terms of the impact on the relevance, we focus on the page-view CTR in practice. The result shows that compared to the relevance-focused ranking, there is only 0.1% drop in the page-view CTR. However, a 3-fold reduction in underdelivery is achieved by our algorithm. Moreover, the number of new items sold in the treatment group is 15.6% more than that in the control group. The results exhibit the effectiveness and practicability of our method.

Additional Related Work

In addition to the literature mentioned in the previous sections, there are many other related work in the generalized traffic shaping family. For example, besides the business requirements considered in this paper, some studies focus on the fairness in ranking (Biega, Gummadi, and Weikum 2018; Singh and Joachims 2018; Geyik, Ambler, and Kenthapadi 2019) or diversity issues (Devanur et al. 2016; Dickerson et al. 2019).

Another important concept in this paper is submodularity, especially the *submodular welfare maximization* (SWM) problem (Vondrák 2008). (Kapralov, Post, and Vondrák 2013) are the first to study the online version of SWM and show that the greedy algorithm is $\frac{1}{2}$-competitive and optimal in the adversarial setting. Recently, the result is improved to 0.5096-competitive under the random permutation assumption by (Korula, Mirrokni, and Zadimoghaddam 2018). The closest work to ours is that of (Esfandiari, Korula, and Mirrokni 2016). They investigate the bi-objective online submodular welfare maximization problem, but some important proofs are missing in their paper.

Conclusions

Traffic shaping has attracted significant attention since the relevance-focused ranking fails to capture many realistic business requirements. We propose a unified framework from the aspect of multi-objective optimization for dealing with traffic shaping problem. The designed algorithm allows us to flexibly adjust the priorities of different objectives. It is robust with any traffic distributions and has rigorous theoretical guarantees. Our framework can naturally incorporate a wide range of application-driven requirements because in practice the agents' utilities always exhibit diminishing returns and multiple objectives are involved. Our work presents many interesting directions for future research. For example, which business requirements can be incorporated in our framework and which can not? Investigating our problem under the stochastic assumption by making use of the past knowledge can further improve the algorithm performance. We ignore the position bias and item mutual influence in this paper, but some studies show that such challenges can be also tackled with the help of submodular optimization (Devanur et al. 2016; Dickerson et al. 2019). We are interested in modeling these all together.

Acknowledgments

This work is supported by Science and Technology Innovation 2030 - "New Generation Artificial Intelligence" Major Project (No. 2018AAA0100902), China. We would like to thank the anonymous reviewers of the paper for comments that provided some interesting food for thought.

References

Agarwal, D.; Chen, B.-C.; Elango, P.; and Wang, X. 2011. Click shaping to optimize multiple objectives. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 132–140.

Agarwal, D.; Chen, B.-C.; Elango, P.; and Wang, X. 2012. Personalized click shaping through lagrangian duality for online recommendation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, 485–494.

Agrawal, S.; Wang, Z.; and Ye, Y. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62(4): 876–890.

Bansal, N.; Gupta, A.; Li, J.; Mestre, J.; Nagarajan, V.; and Rudra, A. 2012. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica* 63(4): 733–762.

Biega, A. J.; Gummadi, K. P.; and Weikum, G. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, 405–414.

Buchbinder, N.; and Naor, J. S. 2007. The Design of Competitive Online Algorithms via a Primal–Dual Approach. *Theoretical Computer Science* 3(2-3): 93–263.

Chen, X. A.; and Wang, Z. 2015. A dynamic learning algorithm for online matching problems with concave returns. *European Journal of Operational Research* 247(2): 379–388.

Devanur, N. R.; and Hayes, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*, 71–78.

Devanur, N. R.; Huang, Z.; Korula, N.; Mirrokni, V. S.; and Yan, Q. 2016. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Transactions on Economics and Computation (TEAC)* 4(3): 1–20.

Devanur, N. R.; Jain, K.; and Kleinberg, R. D. 2013. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, 101–107. SIAM.

Dickerson, J. P.; Sankararaman, K. A.; Srinivasan, A.; and Xu, P. 2019. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1877–1884.

Ding, W.; Govindaraj, D.; and Vishwanathan, S. 2019. Whole Page Optimization with Global Constraints. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3153–3161.

Esfandiari, H.; Korula, N.; and Mirrokni, V. 2016. Bi-objective online matching and submodular allocations. In *Advances in Neural Information Processing Systems*, 2739–2747.

Esfandiari, H.; Korula, N.; and Mirrokni, V. 2018. Allocation with traffic spikes: Mixing adversarial and stochastic models. *ACM Transactions on Economics and Computation (TEAC)* 6(3-4): 1–23.

Geyik, S. C.; Ambler, S.; and Kenthapadi, K. 2019. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2221–2231.

Gupta, R.; Liang, G.; Tseng, H.-P.; Holur Vijay, R. K.; Chen, X.; and Rosales, R. 2016. Email volume optimization at LinkedIn. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 97–106.

Kapralov, M.; Post, I.; and Vondrák, J. 2013. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, 1216–1225. SIAM.

Korula, N.; Mirrokni, V.; and Zadimoghaddam, M. 2018. Online submodular welfare maximization: Greedy beats 1/2 in random order. *SIAM Journal on Computing* 47(3): 1056–1086.

Shah, P.; Soni, A.; and Chevalier, T. 2017. Online Ranking with Constraints: A Primal-Dual Algorithm and Applications to Web Traffic-Shaping. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 405–414.

Singh, A.; and Joachims, T. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2219–2228.

Vondrák, J. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, 67–74.

Zhao, B.; Xu, P.; Shi, Y.; Tong, Y.; Zhou, Z.; and Zeng, Y. 2019. Preference-aware task assignment in on-demand taxi dispatching: An online stable matching approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2245–2252.