

Revenue and User Traffic Maximization in Mobile Short-Video Advertising

Dezhi Ran*
School of Computer Science
Peking University, China
dezhiran@pku.edu.cn

Weiqliang Zheng*
Department of Computer Science
Yale University, US
weiqliang.zheng@yale.edu

Yunqi Li
School of Computer Science
Peking University, China
liyunqi@pku.edu.cn

Kaigui Bian
School of Computer Science, National
Engineering Laboratory for Big Data
Analysis and Applications
Peking University, China
bkg@pku.edu.cn

Jie Zhang
Electronics and Computer Science
University of Southampton, UK
jie.zhang@soton.ac.uk

Xiaotie Deng
Center on Frontiers of Computing
Studies
Peking University, China
xiaotie@pku.edu.cn

ABSTRACT

A new mobile attention economy has emerged with the explosive growth of short-video apps such as TikTok. In this internet market, three types of agents interact with each other: the platform, influencers, and advertisers. A short-video platform encourages its influencers to attract users by creating appealing content through short-form videos and allows advertisers to display their ads in short-form videos. There are two options for the advertisers: one is to bid for platform advert slots in a similar way to search engine auctions; the other is to pay an influencer to make engaging short videos and promote them through the influencer's channel. The second option will generate a higher conversion ratio if advertisers choose the right influencers whose followers match their target market. Although displaying influencer ads will generate less revenue, it is more engaging than platform ads, which is better for maintaining user traffic. Therefore, it is crucial for a platform to balance these factors by establishing a sustainable business agreement with its influencers and advertisers. In this paper, we develop a two-stage solution for a platform to maximize short-term revenue and long-term user traffic maintenance. In the first stage, we estimate the impact of user traffic generated by displaying influencer ads and characterize the user traffic the platform should allocate to influencers for overall revenue maximization. In the second stage, we devise an optimal $(1 - 1/e)$ -competitive algorithm for ad slot allocation. To complement this analysis, we examine the ratio of the revenue generated by our online algorithm to the optimal offline revenue. Our simulation results show that this ratio is 0.94 on average, which is much higher than $(1 - 1/e)$ and outperforms four baseline algorithms.

KEYWORDS

Short-video advertising; revenue maximization; competitive ratio

ACM Reference Format:

Dezhi Ran*, Weiqliang Zheng*, Yunqi Li, Kaigui Bian, Jie Zhang, and Xiaotie Deng. 2022. Revenue and User Traffic Maximization in Mobile Short-Video Advertising. In *Proc. of the 21st International Conference on Autonomous*

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: A demonstration of a series of short videos and the advertising scenario with agents' interactions in a short-video app. Advertisers bid for the platform ad slots or cooperate with influencers to deliver engaging ad videos.

Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAA-MAS, 9 pages.

1 INTRODUCTION

Over the last decade, user-generated content (UGC) platforms such as Facebook and Instagram have made communication through news feeds and video feeds easier. Such platforms have also stimulated the transformation of Internet ad auctions, from classical sponsored links on search engines to banner advertising on Facebook and video advertising on Instagram. The trend of delivering information through videos has never been easier and more mobile. Cisco expected that by 2022, 82% of mobile data would be consumed by videos [7].

Recently, short-video mobile apps have grown tremendously, with TikTok, Vine, and Vigo as typical examples. Short-video feeds,

*These authors contributed equally to this work.

in the form of memes, lip-synced songs, and comedy videos, are succinct, passionate and creative, therefore attractive [32, 37]. We show a series of short videos and the advertising scenario of a typical mobile app that consists of three types of agents: platform, advertisers and influencers in Fig. 1. Individuals can post their clips effortlessly, and many influencers have been more active and successful than ever before [39]. Influencers create interesting and attractive short videos, capturing billions of users who spend considerable time watching these videos [26]. In order to produce engaging content, extensive research has been done to understand app users’ interests and behavior patterns [4–6, 18, 23, 24, 38, 40]. Some users subscribe to influencers and become their followers. Short videos, most lasting for 10 to 15 seconds each, are displayed to users on full screen, one after another. Users can choose to watch them or swipe to the next. Amongst these short videos, usually one in every seven is a video ad. There are two types of ads. One is the *platform ad*, which is offered to advertisers through a platform’s real-time bidding mechanism. The other is *influencer ad*, in which advertisers sponsor their partner influencers to make a personalized ad and display it through the influencers’ channel. In the second case, platforms charge a commission fee, and the ads are delivered to targeted audiences – the influencers’ followers – to increase the conversion rate. So, an advertiser is more likely to reach its target customers quickly through influencer marketing.

In this mobile marketplace, the three types of agents – platform, influencers, and advertisers— interact with each other. Many short-video platforms provide productive influencers with benefits, e.g., allocating extra user traffic, monetary reward, and tailored commercial cooperation. These benefits stimulate active influencers to create high-quality content and attract as many followers as possible, which further accelerates the booming of the platform. The high-quality short videos attract many users who generate billions of impressions per day, which incurs advertising over the platform. Short-video app users often base their content on trends and memes, which can change rapidly as what is popular one week will not necessarily be hot the next. So, advertisers consider marketing through these short-video apps to be highly rewarding – as their content can jump on the trends as soon as possible.

The co-existence of two advertising channels is one of the unique features of short-video platforms. The platform ads enable advertisers to display their ads to any user, but this approach suffers a low conversion rate¹, just like advertising on other user-generated content platforms. The influencer ads have a higher conversion rate, but the ads will mainly be displayed to influencers’ followers. These two advertising approaches complement each other and provide advertisers with more choices. From a platform’s perspective, platform ads generate immediate revenue, but users often do not find them entertaining, so displaying too many platform ads risks a loss of users. The influencer ads generate less revenue than platform ads, as advertisers usually submit a smaller bid price for influencer ads. However, they generate user traffic as well since the embedded ads are intrinsic, natural, and engaging. To influencers, they are happy to play a role in the advert so that they have more chance to be exposed. This way, they get more opportunities to attract

¹ Conversion rates are calculated by taking the number of conversions to a sale dividing the number of total ad displays that can be tracked to a conversion during the same period.

followers. As a return, they have more business collaboration with the platform and advertisers. Therefore, the platform, advertisers, and influencers benefit from this feature.

It is a new challenge for these short-video platforms to balance the short-term revenue generated by platform ads and the long-term impact on user traffic generated by influencer ads. In this paper, we deploy badge design [2] as a tool to maximize a platform’s overall profit that combines short-term revenue and long-term user traffic impact when considering interactions among different types of agents. We tackle this problem by developing a two-stage solution. In the *traffic bonus estimation stage*, the platform estimates the user traffic that would be generated by an influencer ad and determines an optimal impression target. In the *online allocation stage*, we present an optimal online ad slot allocation algorithm that achieves a competitive ratio of $1 - \frac{1}{e}$. The algorithm combines advertisers’ bids and the estimated influencers’ traffic bonus and outputs the best short-video ads allocation that maximizes the platform’s total revenue.

We summarize the contribution of our work as follows.

- We formulate the mobile short-video advertising framework. In this framework, after collecting advertisers’ bids and their ads type and content, a platform needs to allocate ads slots such that its overall revenue, as the sum of advertisers’ payments and long-term user traffic bonus, is maximized.
- We devise an algorithm based on badge design for estimating the user traffic bonus as a result of displaying different influencer ads.
- We devise an online allocation algorithm that aggregates bid prices and user traffic bonus for overall revenue maximization. The algorithm is $(1 - \frac{1}{e})$ -competitive and we show that the bound is tight.
- Going forward, we conduct extensive experiments and show that the devised algorithm outperforms four other baseline algorithms in advertising auctions.

2 RELATED WORK

Badge Design for Participant Incentives. There is a large amount of literature on empirical analysis and models of user behavior in user-generated content platforms. Hamari [17] conducted a two-year field experiment by implementing badges in a service. The experimental results showed that users generally used the service in a significantly more active way. Denny [8] conducted a large-scale randomized, controlled experiment, measuring the impact of incorporating a badge-based achievement system within an online learning tool, and a highly significant positive effect was discovered on the quantity of students’ contributions, without a corresponding reduction in their quality. Drawing evidence from empirical data, Anderson et al. [2] concluded that users indeed value badges and modify their actions to earn badges. They proposed a model of how users behave in response to badges awarded for their actions. A game-theoretical approach was proposed for badge design [10], analyzing the incentives created by two different widely-used badge designs in a model where winning a badge is valued, and the effort is costly.

Sponsored Search Auctions. The sponsored search auctions have been extensively studied in the last 20 years. The Generalized

Second Price model bears a similarity to the Generalized English Auctions and the classical assignment games. In particular, Benjamin *et al.* [11] and Varian [36] discussed their similarities and characterized their equilibria. The seminal work by [28] established the competitive analysis of search engines' revenue maximization problem in sponsored search auctions. Feldman *et al.* [12] studied a free disposal model for online ad assignment where the value of an assignment only includes the highest-weighted impressions assigned to each advertiser. Goel *et al.* [16] introduced an impression-plus-click model and investigated a dominant strategy mechanism design problem when bidders' valuations are consistent. Jeong *et al.* [20] considered the special scenario for advertising in a stream and emphasized the attention level drop of a single user viewing the stream. Rong *et al.* [33] discussed the quantal response equilibrium under bounded rationality assumptions. Gatti *et al.* [13] built a cascade model with contextual externalities and bounded user memory for sponsored search auctions. Abhishek *et al.* [1] designed truthful auctions with multi-arm bandits. Mandal and Narahari [25] proposed an ex-post truthful mechanism for multi-slot sponsored search. Shen *et al.* [34] modelled the bidders' behavior and learnt the set of optimal reserve prices by reinforcement learning. Other surveys on sponsored search auctions can be found in [15], [21], [31], and Chapter 28 in [30].

Video Advertising. With the vigorous growth of video-sharing platforms such as YouTube and social networking services such as Facebook, video advertising has been flourishing [29]. However, research on video ad slots allocation was significantly less than that on the sponsored search auctions. As far as we are aware, Geyik *et al.* [14] was the first to address challenges arising in online video advertising, optimized multiple video-specific performance indicators, including engagement (the percentage of the total duration of a video ad that the users have actively watched) and viewability (the goodness of the ad's location), while subject to the budget constraint. They demonstrated the benefit of the proposed framework via empirical results. Sumita *et al.* [35] took account of the length of a video ad and the time spent watching it by users. Assuming that a user watches allocated video-ads to the end, they simplified the problem and presented an online video-ads allocation algorithm.

3 REVENUE AND USER TRAFFIC MAXIMIZATION

In this section, we formulate the problem of revenue and user traffic maximization in mobile short-video advertising.

On a short-video platform, there is a set of advertisers. Each advertiser i has a fixed daily budget $B(i)$ for real-time bidding (RTB). The platform has a set of advert slots. Each slot j corresponds to a short period in a day. Each advertiser i submits their bid prices $b(i, j)$ for slots j and their budget constraint $B(i)$ to the platform, as well as their short-video ad content. Upon receiving these inputs, the platform must immediately allocate each slot j to an advertiser i^* , and the allocation is not revokable. Therefore, this is an online problem. The platform is interested in maximizing the revenue collected from the advertisers and the user traffic generated by displaying these ads.

An advertiser can submit two types of ads. One is the platform ad which is a short video created by the advertiser that introduces their

products. This type of ad can be displayed to a very broad range of users. The other is the influencer ad where an influencer embeds the advertisement into an engaging short video so that the users enjoy the video content and are more likely to appreciate the products. This type of ad usually is displayed to the influencer's followers only. A unique feature of the influencer ad is that it not only generates revenue for the platform, but also attracts users. Since maintaining as many users as possible is vital to the development of a platform, the platform needs to consider the dual effects of influencer ads and maximize the overall benefit, including the short-term revenue and long-term user traffic bonus. However, a practical consideration for the platform is that displaying any particular advertiser's ads too many times a day may have a negative effect on attracting users, even though this advertiser may have a big budget and high bid prices. Given this structure, we develop a **two-stage solution** for this problem.

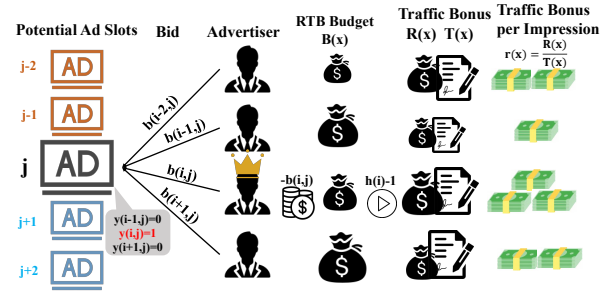


Figure 2: During the online ad slot allocation stage, each advertiser has a budget and traffic impression target. The platform displays ads in real-time.

The Traffic Bonus Estimation Stage. As shown in Fig. 2, the advertisers participate in the online short-video ad slot auction, a.k.a. real-time bidding (RTB). Let n denote the number of advertisers and m the number of ad slots. Assume there are K influencers available for collaboration with advertisers on a platform. With a slight abuse of notation, let i denote the influencer who produces an ad with advertiser i . In case an advertiser i only bids for platform ads, its corresponding influencer i is null. Let $T(i)$ be the impression upper limit that a platform will display advertiser/influencer i 's ads, as going beyond it will bore the users. Let $R(i)$ denote the platform's estimation on the benefit of the user traffic if they display advertiser i 's ad $T(i)$ times.

In the first stage, we devise an algorithm by deploying *badge design* to estimate the bonus of attaining user traffic when displaying influencer ads. Badges or other equivalent rewards are used to recognize a user's contributions to a site. When properly designed, it can be used to gear users' incentive to make an effort to win the badges and hence make a significant contribution to the booming of a platform. The algorithm estimates an upper bound of impressions $T(i)$ the platform will display an advertiser's ad and the user traffic bonus $R(i)$.

The Online Allocation Stage. In the second stage, we devise an online algorithm that balances the bid prices (short-term revenue) and user traffic bonus (long-term impact) to achieve a competitive ratio of $1 - \frac{1}{e}$.

We denote the allocation variable of slot j to the advertiser i by $y(i, j)$. That is, $y(i, j) = 1$ if the advertiser i gets slot j and 0 otherwise. Let $h(i)$ be the outstanding impressions for the advertiser i , and $T(i) - h(i)$ the number of impressions that the advertiser i 's ad has been displayed. Let $r(i) = R(i)/T(i)$ be the estimated user traffic bonus per impression for advertiser i . Given these inputs, the platform's offline optimal revenue is captured by the following linear programming.

$$(\mathcal{P}) \quad \max_{y(i,j), h(i)} \sum_{j=1}^m \sum_{i=1}^n b(i, j)y(i, j) + \sum_{i=1}^n (T(i) - h(i))r(i)$$

$$\text{s.t.} \quad \sum_{i=1}^n y(i, j) \leq 1, \quad \forall j \in \{1, \dots, m\} \quad (1)$$

$$\sum_{j=1}^m b(i, j)y(i, j) \leq B(i), \quad \forall i \in \{1, \dots, n\} \quad (2)$$

$$\sum_{j=1}^m y(i, j) + h(i) \geq T(i), \quad \forall i \in \{1, \dots, n\} \quad (3)$$

$$y(i, j), h(i) \geq 0, \quad \forall i \in \{1, \dots, n\}, j \in \{1, \dots, m\} \quad (4)$$

The first term of the objective function captures the platform's revenue collected from RTB². The second term captures its revenue collected from the user traffic bonus. Constraint (1) guarantees that each slot is allocated to at most one advertiser. Constraint (2) is due to the advertisers' budget limit. Constraint (3) is by the definition of outstanding impressions $h(i)$. When $\sum_{j=1}^m y(i, j) \geq T(i)$, $h(i) = 0$, it indicates the extra traffic allocation is accomplished.

4 ALGORITHMS AND ANALYSIS

In this section, we first propose an algorithm that estimates the user traffic bonus for each influencer. We then present an online ad slot allocation algorithm for short-video advertising and show that it is $(1 - \frac{1}{e})$ -competitive.

4.1 The User Traffic Bonus Estimation Algorithm

Let phase l denote the time span that an influencer i 's short videos are displayed and $N(i, l)$ the number of times they are displayed in phase l . In a phase, some users who viewed influencer i 's short videos will become its follower. For each phase, a platform maintains the *quality level* $q(i, l) \in [0, 1]$, which is a probability estimation of how likely users will become influencer i 's followers in phase l . The higher $q(i, l)$ is, the more users who viewed influencer i 's short videos become i 's follower in phase l . When the platform wants to stimulate the number of influencer i 's followers $N(i, l)q(i, l)$ in phase l , they can do it by increasing the number of times influencer i 's short videos are displayed. Suppose the platform display influencer i 's short videos $T(i, j)$ additional times in this phase. In that case, more users could have become its follower, and influencer i 's quality level would be updated to $p(i, j)$. Therefore, we have that

$$(N(i, l) + T(i, l))q(i, l) = N(i, l)p(i, l).$$

²We adopt the First-Price auction (the bidder with the highest bidding price pays their bid) for two reasons. First, our algorithm and competitive analysis can be generalized to the Second-Price auction while maintaining the same competitive ratio as in the case in [28]. Second, Google recently switched from Generalized Second Price (GSP) to Generalized First Price (GFP) auction in Ad Manager [9, 22]. Although GSP has been used since Google's sponsored search auctions, for various reasons, such as reducing programmatic inefficiencies, it is not dominating GFP everywhere anymore.

Hence, in order to increase influencer i 's followers to $N(i, l)p(i, l)$, the platform needs to allocate influencer i additional user traffic

$$T(i, l) = \frac{p(i, l)N(i, l)}{q(i, l)} - N(i, l). \quad (5)$$

We note that the platform allocates user traffic to an influencer at a cost. The more the influencer i 's short videos are displayed, the less the other influencers' short videos could be displayed as the total traffic in a period is limited. Also, the marginal benefit of allocating more traffic to any particular influencer decreases as users will be bored with viewing the same influencer too many times. Therefore, we assume that the platform maintains a cost function $V(i, x)$ to evaluate the user traffic opportunity cost by allocating x additional user traffic to influencer i . Let $U(i, l)$ denote the platform's utility generated by displaying influencer i 's short videos in phase l . Following [2], we associate the platform's utility with a time discount factor γ . For $\gamma \in [0, 1]$, the discounted utility of the platform is $\sum_{l=0}^{\infty} U(i, l)\gamma^l$.

We denote by $s_l \in \{0, 1, 2, \dots, t(i)\}$ the number of additional followers influencer i attracts through additional user traffic and $U(i, s_l)$ the optimal discounted utility generated by influencer with s_l . Note that $U(i, s_l)$ depends only on s_l , not on the history state. Let $v(i)$ denote the value to the platform that influencer i attracts every additional $t(i)$ followers. We will use $s_l = t(i)$ and $U(i, t(i)) = v(i)$ as the initial input of our algorithm. For $l = 1, 2, \dots, t(i) - 1$, since influencer i can attract a new follower with probability $p(i, l)$. Following [2], the time discount factor γ associates with the probability of reaching the next phase and in each phase there are two potential actions to be taken. With probability $p(i, l)$, a user will become a follower and phase $l + 1$ will be reached, and with probability $1 - p(i, l)$, the user will not follow the influencer. Hence, We have the following equation,

$$U(i, l) = \gamma[p(i, l)U(i, l+1) + (1 - p(i, l))U(i, l)] - V(i, T(i, l)).$$

Thus, we have

$$U(i, l) = \frac{-V(i, \frac{p(i, l)N(i, l)}{q(i, l)} - N(i, l)) + \gamma p(i, l)U(i, l+1)}{1 - \gamma + p(i, l)\gamma},$$

which can be solved efficiently. We can use backward induction to compute $U(i, 0)$ and determine $T(i)$ and $R(i)$ as shown in Algorithm 1.

4.2 The Online Allocation Algorithm

Our insight on designing a competitive online algorithm is to balance the user traffic bonus and bid prices properly. We employ a primal-dual method to construct the least scaling factor to ensure dual feasibility.

Firstly, we derive the dual problem of the primal LP (\mathcal{P}) as follows.

$$(\mathcal{D}) \quad \min_{x(i), \phi(i), z(j)} \sum_{i=1}^n B(i)x(i) + \sum_{i=1}^n T(i)r(i)(1 - \phi(i)) + \sum_{j=1}^m z(j)$$

Algorithm 1 Stage 1: The User Traffic Bonus Estimation Algorithm.

Input: The incremental follower target $t(i)$ and influencer quality level $q(i, l) \forall 0 \leq l \leq t(i)$; original traffic $N(i, l)$; platform value $v(i)$ and platform opportunity cost function $V(i, x)$ for required x .
Output: The additional impression target $T(i)$ and user traffic bonus $R(i)$.

Initialize $T(i, l), p(i, l), T(i), R(i), U(i, l) \leftarrow 0$ for $\forall 0 \leq l \leq t(i)$

- 1: $U(i, t(i)) \leftarrow v(i)$
 - 2: **for** l in $t(i) - 1, \dots, 0$ **do**
 - 3: $p(i, l) \leftarrow \operatorname{argmax}_{p \in [0, 1]} \frac{-V(i, \frac{p(i, l)N(i, l)}{q(i, l)} - N(i, l)) + \gamma p(i, l)U(i, l+1)}{1 - \gamma + p(i, l)\gamma}$
 - 4: $U(i, l) \leftarrow \frac{-V(i, \frac{p(i, l)N(i, l)}{q(i, l)} - N(i, l)) + \gamma p(i, l)U(i, l+1)}{1 - \gamma + p(i, l)\gamma}$
 - 5: $T(i, l) \leftarrow T(i, l+1) + \frac{p(i, l)N(i, l)}{q(i, l)} - N(i, l)$
 - 6: **end for**
 - 7: $T(i) \leftarrow T(i, 0)$
 - 8: $R(i) \leftarrow \max\{U(i, 0), 0\}$
-

$$\text{s.t. } z(j) + b(i, j)x(i) - r(i)\phi(i) \geq b(i, j) \quad (6)$$

$$\forall i \in \{1, \dots, n\}, j \in \{1, \dots, m\}$$

$$0 \leq \phi(i) \leq 1, \quad \forall i \in \{1, \dots, n\} \quad (7)$$

$$x(i), z(j) \geq 0, \quad \forall i \in \{1, \dots, n\}, j \in \{1, \dots, m\} \quad (8)$$

where $z(j)$, $x(i)$, and $\phi(i)$ are dual variables that correspond to constraints (1), (2), and (3) in (\mathcal{P}) , respectively.

Secondly, we will derive an allocation rule that hits two birds with one stone. That is, to ensure dual feasibility and to upper bound the value increase of (\mathcal{D}) by the value increase of (\mathcal{P}) multiplying a constant, so that we can achieve the desired competitive ratio at the same time.

In the following, let $R_{\max} = \max_{i, j} \left\{ \frac{b(i, j)}{B(i)} \right\}$ denote the largest ratio of bid price to budget and $T = \min_i \{T(i)\}$ the minimum extra impression target³. Define the scalar $c = (1 + R_{\max})^{1/R_{\max}}$. In Algorithm 2, the allocation takes into account each advertiser's bid price $b(i, j)$ as well as the bonus-per-impression $r(i)$, which is decided by a weighted sum of these two terms.

We remark that Algorithm 2 balances the short-term revenue and the long-term user traffic bonus for the platform.

We combine the two stages and present Algorithm 3 as below.

4.3 Competitive Ratio Analysis

In this section, we show that the algorithm achieves a competitive ratio $(1 - 1/e)$ under mild assumptions. Firstly, we prove the following proposition which provides an upper bound of the variables $\phi(i)$.

Proposition 1. *During the execution of Algorithm 2, if advertiser i 's impression playback task as per the extra impression target is accomplished, then $\phi(i) < \frac{c}{(c-1)(T(i)+1)}$.*

PROOF. Let $\phi(i)_k$ be the value of $\phi(i)$ after the k^{th} ($k \geq 1$) playback of advertiser i 's ads. In particular, $\phi(i)_0 = 1$. According to the

³To be precise, in the definition of R_{\max} we dismiss the advertisers who do not participate in RTB; in the definition of T , we exclude the advertisers who do not cooperate with influencers or the influencers whose extra impression targets are 0.

Algorithm 2 Stage 2: The Online Allocation Algorithm.

Input: The bid prices $b(i, j)$ and budgets $B(i)$; extra impression target $T(i)$ and bonus-per-impression $r(i)$.

Output: The slot allocation $y(i, j)$ and outstanding impressions $h(i)$.

Initialize $x(i), y(i, j), z(j) \leftarrow 0, \phi(i) \leftarrow 1, h(i) \leftarrow T(i)$.

- 1: Let $i^* = \operatorname{argmax}_i \{b(i, j)(1 - x(i)) + r(i)\phi(i)\}$.
 - 2: Set $y(i^*, j) = 1$. Allocate slot j to advertiser i^* .
 - 3: Charge advertiser i^* by $\min\{b(i^*, j), B(i^*) - \sum_{k < j} b(i^*, k)y(i^*, k)\}$.
 - 4: if $h(i^*) > 0$: $h(i^*) \leftarrow h(i^*) - 1$.
 - 5: Set $z(j) \leftarrow b(i^*, j)(1 - x(i^*)) + r(i^*)\phi(i^*)$.
 - 6: Set $x(i^*) \leftarrow x(i^*) + \frac{b(i^*, j)}{B(i^*)} + \frac{b(i^*, j)}{(c-1)B(i^*)}$.
 - 7: Set $\phi(i^*) = \max\{0, \phi(i^*)(1 + \frac{1}{T(i^*)}) - \frac{c}{(c-1)T(i^*)}\}$
 - 8: if $h(i^*) = 0$: Set $\phi(i^*) = 0$
-

Algorithm 3 The combined Algorithm.

Input: The influencer quality level $q(i, l)$ and incremental follower target $t(i)$; original traffic $N(i, l)$; platform value $v(i)$ and platform opportunity cost function $V(i, x)$. The bid prices $b(i, j)$ and budgets $B(i)$ for $\forall 1 \leq i \leq n, 1 \leq j \leq m, 0 \leq l \leq t(i)$ and required x .
Output: The slot allocation $y(i, j)$ and outstanding impressions $h(i)$.

- 1: **for** i in $1, \dots, n$ **do**
 - 2: **Compute** $T(i), R(i)$ using the user traffic bonus estimation algorithm:
 $T(i), R(i, j) \leftarrow \text{Algorithm 1}(q(i, l), t(i), N(i, l), v(i), V(i, x))$.
 - 3: $r(i) \leftarrow \frac{R(i)}{T(i)}$
 - 4: **end for**
 - 5: **for** j in $1, \dots, m$ **do**
 - 6: **Compute** $y(i, j), h(i)$ using the online allocation algorithm:
 $y(i, j), h(i) \leftarrow \text{Algorithm 2}(b(i, j), B(i), T(i), r(i))$
 - 7: **end for**
-

update of $\phi(i)$ during the execution of the algorithm (line 7), the closed-form of $\phi(i)_k$ can be written as

$$\phi(i)_k = \begin{cases} \max\{0, \frac{c - (1 + \frac{1}{T(i)})^k}{c-1}\} & \text{if } \phi(i)_{k-1} > 0 \\ 0 & \text{if } \phi(i)_{k-1} = 0 \end{cases}$$

If $\phi(i)_{T(i)} = 0$, the proposition holds. If $\phi(i)_{T(i)} > 0$, then one more iteration leads to

$$\phi(i)_{T(i)} \left(1 + \frac{1}{T(i)}\right) - \frac{c}{(c-1)T(i)} = \frac{c - (1 + \frac{1}{T(i)})^{T(i)+1}}{c-1}$$

For $T(i) \geq 1$, it holds that

$$c = (1 + R_{\max})^{\frac{1}{R_{\max}}} < e < \left(1 + \frac{1}{T(i)}\right)^{T(i)+1},$$

so, $\phi(i)_{T(i)} \left(1 + \frac{1}{T(i)}\right) - \frac{c}{(c-1)T(i)} < 0$, indicating $\phi(i)_{T(i)} < \frac{c}{(c-1)(T(i)+1)}$, which completes the proof. \square

Then we show the main theorem.

Theorem 1. *The competitive ratio of Algorithm 2 is $1 - \frac{1}{e}$.*

PROOF. Let P and D denote the objective functions of the primal problem (\mathcal{P}) and the dual problem (\mathcal{D}), respectively. We prove the theorem by showing three facts:

(i) At the completion of allocating m ad slots, the value of $x(i)$, $\phi(i)$, and $z(j)$ is a feasible solution to the dual problem;

(ii) $D \leq (1 + \frac{1}{c-1})(1 + \frac{c}{(c-1)(T+1)}) \cdot P$;

(iii) The value of $y(i, j)$ and $h(i)$ is an almost feasible solution to the primal problem, with slightly violating constraint (2).

Proof of (i). For any slot j , according to the allocation rule defined by line 1 and 2 of the algorithm and the update of $z(j)$ defined by line 5, the Constraint (5) of (\mathcal{D}) is satisfied for every advertiser. In addition, $x(i)$, $\phi(i)$, and $z(j)$ are non-negative; $\phi(i)$ is initialized to be 1 and is non-increasing according to line 7. So, the solution is feasible for the dual problem.

Proof of (ii). Define $P = P_1 + P_2$, where $P_1 = \sum_{j=1}^m \sum_{i=1}^n b(i, j)y(i, j)$ and $P_2 = \sum_{i=1}^n (T(i) - h(i))r(i)$. For every slot j , according to the line 5 of the algorithm, define $z(j) = z_1(j) + z_2(j)$, where $z_1(j) = b(i^*, j)(1 - x(i^*))$ and $z_2(j) = r(i^*)\phi(i^*)$. Similarly, define $D = D_1 + D_2$, where $D_1 = \sum_{i=1}^n B(i)x(i) + \sum_{j=1}^m z_1(j)$ and $D_2 = \sum_{i=1}^n T(i)r(i)(1 - \phi(i)) + \sum_{j=1}^m z_2(j)$.

When allocating slot j , let D_1^j and P_1^j denote the increase of D_1 and P_1 during the j -th iteration. Obviously, $D_1^j = (1 + \frac{1}{c-1})b(i, j) = (1 + \frac{1}{c-1})P_1^j$. Since it holds for every iteration, summing over $j = 1, \dots, m$ leads to the equality

$$D_1 = \left(1 + \frac{1}{c-1}\right)P_1. \quad (9)$$

Next, we show that

$$D_2 \leq \left(1 + \frac{1}{c-1}\right) \left(1 + \frac{c}{(c-1)(T+1)}\right) P_2. \quad (10)$$

Let $P_{2,i}^j$ and $D_{2,i}^j$ denote the value of P_2 and D_2 derived from allocating slot j to advertiser i , respectively. Let $P_{2,i} = \sum_j P_{2,i}^j$ and $D_{2,i} = \sum_j D_{2,i}^j$. At the j -th iteration when the algorithm is allocating the slot, there are three possible cases according to the extend to which advertiser i 's contract has accomplished.

If $h(i) = 0$, i.e., advertiser i 's (influencer i 's) extra impression target is already accomplished, $h(i)$ and $\phi(i)$ are equal to 0 and won't change anymore. So, we have that

$$D_{2,i}^j = P_{2,i}^j = 0.$$

If $h(i) > 1$, i.e., advertiser i 's (influencer i 's) extra impression target has multiple outstanding impressions, then

$$D_{2,i}^j \leq \left(1 + \frac{1}{c-1}\right)r(i) = \left(1 + \frac{1}{c-1}\right)P_{2,i}^j \quad (11)$$

If $h(i) = 1$, i.e., advertiser i 's (influencer i 's) extra impression target has only one outstanding impression, (11) also holds. However, the algorithm would set $\phi(i) = 0$, which may cause an extra increase $\text{ext}(i)$. According to Proposition 1, we have $\text{ext}(i) = T(i)r(i)\phi(i)_{T(i)} < T(i)r(i) \cdot \frac{c}{(c-1)(T(i)+1)} = \left(\frac{cT(i)}{(c-1)(T(i)+1)}\right)r(i)$. Should advertiser i 's (influencer i 's) extra impression target be accomplished with the allocation of slot j , we have that $D_{2,i} =$

$T(i)r(i)$. Therefore, $\text{ext}(i)$ will not cause a large deviation from $D_{2,i}$. In fact, it holds that

$$\text{ext}(i) + D_{2,i} < \left(1 + \frac{c}{(c-1)(T(i)+1)}\right)D_{2,i}. \quad (12)$$

As (12) holds for every advertiser i , we have that

$$\begin{aligned} D_2 &\leq \sum_i \text{ext}(i) + D_{2,i} \\ &< \sum_i \left(1 + \frac{c}{(c-1)(T(i)+1)}\right)D_{2,i} \\ &\leq \sum_i \left(1 + \frac{1}{c-1}\right) \left(1 + \frac{c}{(c-1)(T+1)}\right)P_{2,i} \\ &= \left(1 + \frac{1}{c-1}\right) \left(1 + \frac{c}{(c-1)(T+1)}\right)P_2 \end{aligned}$$

In all cases, (10) holds. The fact (ii) follows immediately by combining (9) and (10).

Proof of (iii). It is obvious that the algorithm execution always meets the constraints (1), (3), and (4). The only slight violation to the budget constraint (2) may happen when an advertiser wins an ad slot as per line 2 of the algorithm, but its outstanding budget is less than $b(i, j)$. Following a similar proof to [3], we can conclude that when $\sum_j b(i, j)y(i, j) \geq B(i)$, $x(i) \geq 1$ holds for $i \in \{1, \dots, n\}$. Therefore, there can be at most one iteration in which the advertiser is charged less than $b(i, j)$. Hence, it holds that

$$\sum_j b(i, j)y(i, j) \leq B(i) + \max_j b(i, j).$$

So, (iii) is true.

Let OPT denote the optimal offline revenue and PR the online revenue achieved by Algorithm 2. Let $PR = PR_1 + PR_2$, where $PR_1 = \sum_i PR_{1,i} = \sum_i \sum_j b(i, j)y(i, j)$ is the revenue collected from RTB and R_2 the user traffic bonus collected from the extra impression allocation. By (iii), for each advertiser i ,

$$PR_{1,i} \geq P_{1,i} \cdot \frac{B(i)}{B(i) + \max_j b(i, j)} \geq P_{1,i} \cdot (1 - R_{\max}).$$

Summarize over i we get $PR_1 \geq (1 - R_{\max})P_1$. Also, $PR_{2,i} = P_{2,i}$, $\forall i$ implies $PR_2 = P_2$. Therefore,

$$\begin{aligned} PR &= PR_1 + PR_2 \geq (1 - R_{\max})P_1 + P_2 \\ &\geq (1 - R_{\max})(P_1 + P_2) = (1 - R_{\max})P \\ &\geq (1 - R_{\max}) \left(1 - \frac{1}{c}\right) \left(1 - \frac{1}{\frac{c-1}{c}(T+1) + 1}\right) \cdot D \\ &\geq (1 - R_{\max}) \left(1 - \frac{1}{c}\right) \left(1 - \frac{1}{\frac{c-1}{c}(T+1) + 1}\right) \cdot OPT \end{aligned}$$

where the second to last inequality is due to (ii) and the last equality due to the principle of weak duality. This gives us the competitive ratio of $(1 - \frac{1}{c})(1 - \frac{1}{\frac{c-1}{c}(T+1)+1})(1 - R_{\max})$. Following the literature, a bid price is often considered much smaller than one's budget. So, R_{\max} tends to 0. Therefore, $c = (1 + R_{\max})^{1/R_{\max}}$ approaches e . Furthermore, the impression tasks T are usually large. Given these, we conclude that the competitive ratio is $1 - \frac{1}{e}$. \square




Figure 3: The Revenue Algorithm3-to-OPT ratio on average.




Figure 4: Revenue per slot of compared algorithms.




Figure 5: Revenue per slot vs. different proportions of cooperative advertisers.

In addition, the ratio is tight.

Corollary 1. *The competitive ratio $1 - 1/e$ is tight.*

The proof follows the fact that our short-video advertising framework degenerated to Adwords problem [28] when the user traffic bonus $R(i) = 0$ and all advertisers’ budget are used on RTB. A lower bound of $1 - 1/e$ is proved for randomized algorithms for Adwords problem [28]. We remark that Algorithm 2 is time-efficient. It selects the advertiser i^* returned by Algorithm 2 and updates the variables of the advertiser i^* . Thus, selecting the maximum dominates the run-time of the algorithm, i.e., $O(nm)$ in total.

5 EXPERIMENTAL RESULTS

5.1 Experimental Setting

Data acquisition and processing. We construct an ad auction dataset from a short-video playback dataset [19]. Each piece of data in the dataset records the video playback history information, including user id, video id, and the frequency and duration that users have watched the videos. We select the top 1,000 most-frequently

played videos as short-video ads, and regard each of them as owned by an individual advertiser. We further divide these 1,000 advertisers into two categories: 50% of the advertisers will cooperate with influencers to bid for influencer ad (*cooperative advertisers* for short) and 50% of the advertisers will only bid for platform ad (*non-cooperative advertisers* for short). Given the dataset [19], we observe that the density curves of $R(i)$ and $T(i)$ exhibit a bell shape, so we assume they follow Gaussian distributions. We compute their means and use the 68–95–99.7 rule to find the variance best fits the data points. To conclude, $R(i)$ and $T(i)$ identically and independently follow the Gaussian distribution $N(15, 1)$ and $N(120, 49)$, respectively. The advertisers’ budgets $B(i)$ are independent and identically drawn from the Gaussian distribution $N(150, 64)$. For influencers, their quality levels $q(i)$ are i.i.d. variables drawn from the Gaussian distribution $N(0.5, 0.01)$. Based on the observation that the ads that are played more often would better fit the content of the short videos, we set the bid prices $b(i, j)$ proportional to the average watching duration and playback times. We set the number of users and the number of ad slots as 5,000 and 250,000, respectively. The code of the experiments is available online here.⁴

Baseline algorithms. For a comprehensive comparison, apart from our algorithm, we implement four other algorithms. One of them is a direct adaptation of the algorithm devised by Buchbinder et al. [3], and the other three employ heuristic strategies in the short-video advertising framework.

Budget-Constrained (BC) Algorithm [3]: The algorithm is a basic primal-dual online one devised in the classical ad auction context with a tight competitive ratio of $1-1/e$. To implement it on our dataset, we keep the advertiser i ’s budget $B(i)$ and their bid price $b(i, j)$ the same as ours, while ignoring the user traffic bonus.

Traffic Bonus-First (TBF) Algorithm: This algorithm takes extra impression target completion as its priority, so its main difference to our algorithm is the set of advertisers that is taken into consideration in Algorithm 2, Line 1. While the argmax operation of our algorithm takes all advertisers into consideration, TBF selects the most suitable advertiser in the set of advertisers who still have unaccomplished impression targets to win the ad slot. After all extra impression targets are completed, TBF performs the same as our algorithm.

Bid-Price-First (BPF) Algorithm [27]: This is a greedy algorithm based on the first-price sealed-bid auction that directly chooses the advertiser who is willing to pay the highest bid price and is available to pay that amount for the slot, with a competitive ratio of $1/2$.

Virtual-Bonus (VB) Algorithm: This algorithm is an adaptation of the BC algorithm by assuming that the platform has an extra impression target $T(i)$ for each advertiser (influencer), but there is no option to reach a consensus with the platform explicitly. In this context, the platform gets $b(i, j) + r(i)$ when the advertiser i have not got $T(i)$ slots yet and $b(i, j)$ once the advertiser i have achieved $T(i)$ extra impressions. The allocation and pricing of the algorithm are the same as the BC algorithm.

⁴https://github.com/TreceyJueves/Revenue_traffic_maximize_Al




Figure 6: The average revenue gained by advertisers.




Figure 7: The percentage of slots allocated to different types of advertisers.




Figure 8: The impression target-completion rate (TCR) of influencers.

5.2 Results

Revenue Algorithm3-to-OPT ratio. While the competitive analysis provides us the worst-case guarantee of the revenue achieved by our algorithm, compared to the optimal (OPT) revenue, we are interested in how much the revenue ratio is on average when the parameters vary. We randomly choose 10 to 30 advertisers from the set of advertisers and sample 2,000 ad slots. Fig. 3 shows that the revenue ratio is over 0.93, much better than $(1 - \frac{1}{e}) \approx 0.632$.

Revenue per slot of compared algorithms. Fig. 4 shows the revenue comparison of our algorithm against the other baseline algorithms. We construct a video playback trace set with 5,000-50,000 slots by sampling the dataset uniformly at random and run the experiments 1,000 times. We construct 100 different bidding scenarios with the same average bidding price per slot such that the experiment settings are more comprehensive. In some situations, the slots available for bidding run up before advertisers’ budgets while others do not. Since slots are usually in short supply for the platform and are the only resource it has, the platform needs to observe the average revenue per slot. Compared to other algorithms, Algorithm 3 exhibits the greatest superiority by generating the highest revenue per slot for the platform. The VB algorithm, which to

some extent, resembles our algorithm, achieves the second-highest average revenue. Results show that taking user traffic bonus into consideration indeed improves the platform’s revenue. **Sensitivity analysis.** Fig. 5 shows the influence of altering the proportion of cooperative advertisers. The results show that as the percentage of advertisers cooperating with influencers increases, the revenue of the platform gained from each slot increases as well. Our algorithm achieves the best performance among all algorithms by providing the platform with the highest revenue all along.

Incentive for advertiser participation. Fig. 6 shows the advertising profit gained by advertisers using different algorithms throughout the experiment. It is generally recognized that influencer ads have higher conversion rates than platform ads, so we set the conversion rate of cooperative advertisers 10% higher than non-cooperative advertisers. It can be inferred from our experiment results that during the first one-third of an auction, our algorithm is among the three leading algorithms in average advertiser revenue per slot and becomes the sole leader in the last two-thirds. This provides a great incentive for advertisers to accept a bidding system using our algorithm. Fig. 7 shows that under different conditions, advertisers who cooperate with influencers can always get more slots for advertisement than those who do not cooperate, by approximately 8%, while the numbers of the two types of advertisers are the same. Combining with the higher conversion rate of influencer ads, it incentivizes advertisers to cooperate with influencers.

Impression target completion rate (TCR). Fig. 8 show the extra impression target completion rate of influencers. The extra impression target completion rate (TCR) is the proportion of influencers who have completed their targets. Normally, besides finishing targets, influencers would like to display their videos smoothly in a way that spans a reasonable period, rather than having users consume all their impressions in a short time. That is to say, different influencers expect to complete their targets as close as possible, and thus the ideal state of TCR curve would be a rapid rise in the short term, indicating the fairness amongst influencers. As shown in figure 8, Traffic Bonus-First (TBF) algorithm is the benchmark for TCR, whose TCR curve is the steepest. The TCR curves of BPF and VB are relatively flat, indicating the undesired impression target completion scheme. Algorithm 3 and BC perform similar to TBF, which are much better in fulfilling influencers expectations.

6 CONCLUSION

With the recent rise of short-video apps and the reforming short-video ads allocation problem, in this paper, we devise Algorithm 3 consisting of a user traffic bonus estimation algorithm and a $(1-1/e)$ -competitive algorithm that maximize the short-term revenue and long-term user traffic of the platform. Experimental results show that the revenue achieved by the framework against the optimal revenue is much higher than the worst-case ratio of 0.632, and it achieves a higher revenue than the other four baseline methods.

ACKNOWLEDGEMENT

This work is partially supported by National Key R&D Program of China (2020AAA0105200), Beijing Academy of Artificial Intelligence (BAAI), NSFC 62032003 and 61632017, BJNSF L192004, and a Leverhulme Trust Research Project Grant (2021-2024).

REFERENCES

- [1] Kumar Abhishek, Shweta Jain, and Sujit Gujar. 2020. Designing Truthful Contextual Multi-Armed Bandits based Sponsored Search Auctions. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 1732–1734.
- [2] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. 2013. Steering user behavior with badges. In *Proceedings of the 22nd international conference on World Wide Web*. 95–106.
- [3] Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. 2007. Online Primal-Dual Algorithms for Maximizing Ad-Auctions Revenue. In *Algorithms – ESA 2007*, Lars Arge, Michael Hoffmann, and Emo Welzl (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 253–264.
- [4] Jingyuan Chen, Xuemeng Song, Liqiang Nie, Xiang Wang, Hanwang Zhang, and Tat-Seng Chua. 2016. Micro tells macro: predicting the popularity of micro-videos via a transductive model. In *Proceedings of the 24th ACM international conference on Multimedia*. 898–907.
- [5] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. 2017. Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. In *SIGIR*. 335–344.
- [6] Xusong Chen, Dong Liu, Zheng-Jun Zha, Wengang Zhou, Zhiwei Xiong, and Yan Li. 2018. Temporal Hierarchical Attention at Category- and Item-Level for Micro-Video Click-Through Prediction. In *MM*.
- [7] VNI Cisco. 2019. Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper. *Porto Salvo, Lisboa. Disponível em:* <https://www.cisco.com/c/pt_pt/about/press/news-archive-2018/20181127.html>, Acesso em 17 (2019).
- [8] Paul Denny. 2013. The effect of virtual achievements on student engagement. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 763–772.
- [9] Stylianos Despotakis, R Ravi, and Amin Sayedi. 2019. First-price Auctions in Online Display Advertising. Available at SSRN 3485410 (2019).
- [10] David Easley and Arpita Ghosh. 2016. Incentives, gamification, and game theory: an economic approach to badge design. *ACM Transactions on Economics and Computation (TEAC)* 4, 3 (2016), 1–26.
- [11] Benjamin G. Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97, 1 (2007), 242–259.
- [12] Jon Feldman, Nitish Korula, Vahab S. Mirrokni, S. Muthukrishnan, and Martin Pál. 2009. Online Ad Assignment with Free Disposal. In *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*. 374–385.
- [13] Nicola Gatti, Marco Rocco, Paolo Serafino, and Carmine Ventre. 2015. Cascade Model with Contextual Externalities and Bounded User Memory for Sponsored Search Auctions. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind (Eds.). ACM, 1831–1832. <http://dl.acm.org/citation.cfm?id=2773459>
- [14] Sahin Cem Geyik, Sergey Faleev, Jianqiang Shen, Sean O'Donnell, and Santanu Kolay. 2016. Joint optimization of multiple performance metrics in online video advertising. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 471–480.
- [15] Zhabiz Gharibshah, Xingquan Zhu, Arthur Hainline, and Michael Conway. 2020. Deep learning for user interest and response prediction in online display advertising. *Data Science and Engineering* 5, 1 (2020), 12–26.
- [16] Sharad Goel, Sébastien Lahaie, and Sergei Vassilvitskii. 2009. Contract auctions for sponsored search. In *International Workshop on Internet and Network Economics*. Springer, 196–207.
- [17] Juho Hamari. 2017. Do badges increase user activity? A field experiment on the effects of gamification. *Computers in human behavior* 71 (2017), 469–478.
- [18] Jun He, Hongyan Liu, Yiqing Zheng, Shu Tang, Wei He, and Xiaoyong Du. 2020. Bi-labeled LDA: Inferring interest tags for non-famous users in social network. *Data Science and Engineering* 5, 1 (2020), 27–47.
- [19] ICMChallenge. 2019. SHORT VIDEO UNDERSTANDING CHALLENGE. <http://ai-lab-challenge.bytedance.com/tce/vc/>.
- [20] Samuel Ieong, Mohammad Mahdian, and Sergei Vassilvitskii. 2014. Advertising in a stream. In *Proceedings of the 23rd international conference on World wide web*. ACM, 29–38.
- [21] Bernard J. Jansen and Tracy Mullen. 2008. Sponsored search: an overview of the concept, history, and technology. *IJEB* 6, 2 (2008), 114–131. <https://doi.org/10.1504/IJEB.2008.018068>
- [22] Renato Paes Leme, Balasubramanian Sivan, and Yifeng Teng. 2020. Why Do Competitive Markets Converge to First-Price Auctions? *CoRR* abs/2001.06074 (2020). [arXiv:2001.06074](https://arxiv.org/abs/2001.06074) <https://arxiv.org/abs/2001.06074>
- [23] Meng Liu, Liqiang Nie, Meng Wang, and Baoquan Chen. 2017. Towards Micro-Video Understanding by Joint Sequential-Sparse Modeling. In *Proceedings of the 25th ACM International Conference on Multimedia (Mountain View, California, USA) (MM '17)*. Association for Computing Machinery, New York, NY, USA, 970–978. <https://doi.org/10.1145/3123266.3123341>
- [24] Jingwei Ma, Guang Li, Mingyang Zhong, Xin Zhao, Lei Zhu, and Xue Li. 2018. LGA: latent genre aware micro-video recommendation on social media. *Multimedia Tools and Applications* 77, 3 (2018), 2991–3008.
- [25] Debmalya Mandal and Yadati Narahari. 2014. A novel ex-post truthful mechanism for multi-slot sponsored search auctions. In *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14, Paris, France, May 5-9, 2014*, Ana L. C. Bazzan, Michael N. Huhns, Alessio Lomuscio, and Paul Scerri (Eds.). IFAAMAS/ACM, 1555–1556. <http://dl.acm.org/citation.cfm?id=2616059>
- [26] Mary Meeker and Liang Wu. 2018. Internet trends report 2018. *Kleiner Perkins*, May 30 (2018), 2018.
- [27] A. Mehta. 2013. *Online Matching and Ad Allocation*.
- [28] Aranyk Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. AdWords and Generalized Online Matching. *J. ACM* 54, 5, Article 22 (Oct. 2007). <https://doi.org/10.1145/1284320.1284321>
- [29] Tao Mei, Xian-Sheng Hua, Linjun Yang, and Shipeng Li. 2007. VideoSense: towards effective online video advertising. In *Proceedings of the 15th ACM international conference on Multimedia*. 1075–1084.
- [30] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511800481>
- [31] Tao Qin, Wei Chen, and Tie-Yan Liu. 2014. Sponsored Search Auctions: Recent Advances and Future Directions. *ACM TIST* 5, 4 (2014), 60:1–60:34. <https://doi.org/10.1145/2668108>
- [32] Miriam Redi, Neil O'Hare, Rossano Schifanella, Michele Trevisiol, and Alejandro Jaimes. 2014. 6 seconds of sound and vision: Creativity in micro-videos. In *CVPR*. 4272–4279.
- [33] Jiang Rong, Tao Qin, and Bo An. 2015. Computing Quantal Response Equilibrium for Sponsored Search Auctions. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind (Eds.). ACM, 1803–1804. <http://dl.acm.org/citation.cfm?id=2773445>
- [34] Weiran Shen, Binghui Peng, Hanpeng Liu, Michael Zhang, Ruohan Qian, Yan Hong, Zhi Guo, Zongyao Ding, Pengjun Lu, and Pingzhong Tang. 2020. Reinforcement Mechanism Design: With Applications to Dynamic Pricing in Sponsored Search Auctions. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, New York, NY, USA, February 7-12, 2020*. AAAI Press, 2236–2243.
- [35] Hanna Sumita, Yasushi Kawase, Sumio Fujita, Takuro Fukunaga, and RIKEN AIP Center. 2017. Online Optimization of Video-Ad Allocation.. In *IJCAI*. 423–429.
- [36] Hal Varian. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6 (2007), 1163–1178. <https://EconPapers.repec.org/RePEc:eee:indorg:v:25:y:2007:i:6:p:1163-1178>
- [37] Yunwen Wang. 2020. Humor and camera view on mobile short-form video apps influence user experience and technology-adoption intent, an example of TikTok (DouYin). *Computers in Human Behavior* 110 (2020).
- [38] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal Graph Convolution Network for Personalized Recommendation of Micro-video. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1437–1445.
- [39] Yanzen Yue, Xiaojuan Ma, and Zhenhui Jiang. 2016. Influence of content layout and motivation on users' herd behavior in social discovery. In *SIGCHI*. 5715–5719.
- [40] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep reinforcement learning for page-wise recommendations. In *RecSys*. 95–103.