

Optimizing Trade-offs Among Stakeholders in Real-Time Bidding by Incorporating Multimedia Metrics

Xiang Chen
National University of Singapore
chxiang@comp.nus.edu.sg

Bowei Chen
University of Lincoln
bchen@lincoln.ac.uk

Mohan Kankanhalli
National University of Singapore
mohan@comp.nus.edu.sg

ABSTRACT

Displaying banner advertisements (in short, ads) on webpages has usually been discussed as an Internet economics topic where a publisher uses auction models to sell an online user's page view to advertisers and the one with the highest bid can have her ad displayed to the user. This is also called *real-time bidding* (RTB) and the ad displaying process ensures that the publisher's benefit is maximized or there is an equilibrium in ad auctions. However, the benefits of the other two stakeholders – the advertiser and the user – have been rarely discussed. In this paper, we propose a two-stage computational framework that selects a banner ad based on the optimized trade-offs among all stakeholders. The first stage is still auction based and the second stage re-ranks ads by considering the benefits of all stakeholders. Our metric variables are: the publisher's revenue, the advertiser's utility, the ad memorability, the ad click-through rate (CTR), the contextual relevance, and the visual saliency. To the best of our knowledge, this is the first work that optimizes trade-offs among all stakeholders in RTB by incorporating multimedia metrics. An algorithm is also proposed to determine the optimal weights of the metric variables. We use both ad auction datasets and multimedia datasets to validate the proposed framework. Our experimental results show that the publisher can significantly improve the other stakeholders' benefits by slightly reducing her revenue in the short-term. In the long run, advertisers and users will be more engaged, the increased demand of advertising and the increased supply of page views can then boost the publisher's revenue.

KEYWORDS

Display advertising; real-time bidding; ad recommendation; trade-offs optimization

ACM Reference format:

Xiang Chen, Bowei Chen, and Mohan Kankanhalli. 2017. Optimizing Trade-offs Among Stakeholders in Real-Time Bidding by Incorporating Multimedia Metrics. In *Proceedings of SIGIR'17, August 7–11, 2017, Shinjuku, Tokyo, Japan*, 10 pages.
DOI: <http://dx.doi.org/10.1145/3077136.3080802>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR'17, August 7–11, 2017, Shinjuku, Tokyo, Japan

© 2017 ACM. 978-1-4503-5022-8/17/08...\$15.00

DOI: <http://dx.doi.org/10.1145/3077136.3080802>

1 INTRODUCTION

Over the last few years, a significant development in display advertising is the emergence of RTB. It is a non-guaranteed delivery system in which page views (also called *impressions*) are bought and sold via programmatic instantaneous auctions [15]. The success of RTB can be largely attributed to its user targeting and wide access. In RTB, advertisers bid for an impression from their targeted user and this impression can be from any webpage from any ad network or platform. If an advertiser wins the auction, her ad will be selected and be displayed to the user in almost real time.

RTB has two major limitations. The first limitation is that it can not provide guaranteed delivery to advertisers. Therefore, those who want to secure and access future impressions in advance will not be satisfied [29]. The solution is to develop a new system that enables time-dependent allocation and pricing of future impressions. Notable examples include the recently discussed programmatic guarantee and ad options [8–11, 41]. The second limitation of RTB is that those displayed ads may not fit their hosting webpages well due to user targeting. It can be: the ad content is irrelevant to the webpage content; the color of the ad image is too light or dark to sufficiently contrast with the webpage color scheme. Ill-fitting ads will affect the effectiveness of ad branding for the advertiser, or even worse annoy users, and potentially reduce the publisher's webpage visits and revenue. It should be noted that this limitation only exists in display advertising, since sponsored search ads are displayed in terms of textual links where visual effects are not very important and also these ads are keyword-based which are well matched with the user's search query. In display advertising, this limitation is worth investigating. Firstly, it has received surprisingly little attention in the recent studies. Secondly, it is an interesting topic that lies in the area combining the research of marketing, multimedia, recommendation, and Internet economics.

In this paper, we discuss a solution to the second limitation of RTB by proposing a computational framework that optimizes trade-offs among stakeholders. As shown in Fig. 1, the framework has two stages. The first stage is based on the existing ad auction model used in RTB; and the second stage re-ranks ads based on the optimal trade-offs. In the re-ranking, there are six metric variables: the publisher's revenue, the advertiser's utility, the ad memorability, the ad CTR, the contextual relevance and the visual saliency. Revenue is always the key concern for the publisher. From an advertiser's perspective, her utility is the short-term benefit and the ad memorability is the long-term benefit. The rest three metric variables represent the user's benefits. Under the proposed framework, we discuss an algorithm that determines the optimal weights of the metric variables. In the paper, we validate the proposed framework on several datasets, including ad auction datasets and multimedia datasets. The ground truth and possible weights combinations for

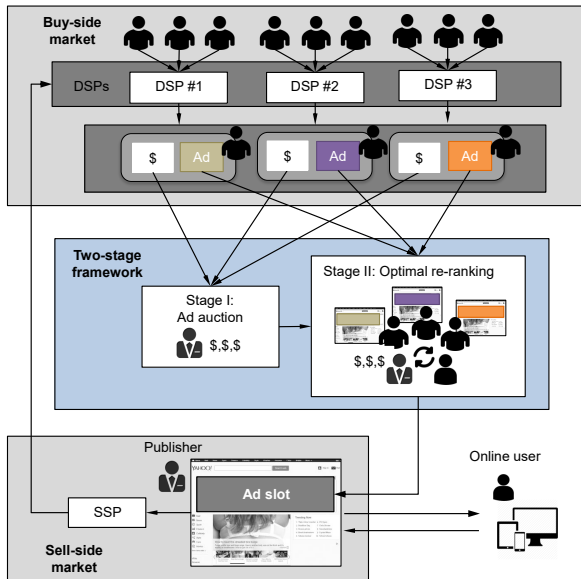


Figure 1: Schematic view of the two-stage framework.

re-ranking are discussed. Our experimental results show that a slight decrease in the publisher’s revenue can increase the performance of other variables significantly in the short term, which will also improve the engagements of both advertisers and users in the long term, and further increase the publisher’s revenue.

This paper has two major contributions. Firstly, the user’s benefits have been taken into account. This is similar to the work of Bachrach et al. [3]. Compared to their work, we increase the dimension of benefits by discussing trade-offs among stakeholders’ multiple benefits. The framework is also tailored to RTB rather than sponsored search. Our proposed model can be regarded as a generalized framework that non-guaranteedly delivers banner ads in RTB. Secondly, multimedia metrics have been introduced to measure some stakeholders’ benefits, such as the ad memorability, the contextual relevance, and the visual saliency. These metrics naturally fit display banner ads but have not yet been discussed along with auction models in previous studies. To the best of our knowledge, this is the first work that combines multimedia metrics and auction theory together.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 discusses the proposed two-stage framework. Section 4 presents our experimental results, and Section 5 concludes the paper.

2 RELATED WORK

In online advertising, the generalized second-price (GSP) auction [14, 40] and the Vickrey-Clark-Groves (VCG) auction [32] have been widely used on different platforms, including display and search domains. Revenue maximization is always the key issue. Lahaie and Pennock [19] proposed to use a squashing parameter to increase revenue of GSP auctions in sponsored search, and then Lahaie and McAfee [18] showed that it can improve social welfare as well. Ostrovsky and Schwarz [30] investigated the effects of optimal reserve prices for GSP auctions on Yahoo’s platform. Thompson and

Leyton-Brown [39] studied a variety of ways of increasing revenue of GSP auctions, including reserve prices and squashing parameter. Furthermore, Yuan et al. [42] implemented a large scale experiment and empirically compared algorithms of setting reserve price for maximizing revenue in RTB.

Trade-offs among multiple objectives or stakeholders have been discussed in several works. Likhodedov and Sandholm [23] proposed a framework that linearly combines revenue and social welfare in a single-item auction. Radlinski et al. [34] investigated the optimization of relevance and revenue in sponsored search. Similarly, Liao et al. [22] combined revenue and relevance for video advertising. Lucier et al. [24] discussed the revenue of GSP auction at equilibrium, and demonstrated that revenue can be maximized at a non-envy-free equilibrium that also generates a socially inefficient allocation. Roberts et al. [36] discussed a new ranking algorithm for sponsored search auction that uses reserve prices to order the ads and discussed conditions under which revenue can be increased. Bachrach et al. [3] proposed a framework that linearly combines objectives relevant to the search engine (revenue), the advertiser (welfare), and the user (clicks).

Most related advertising research focuses on the economic benefits, little attention has been paid to the advertiser’s branding effectiveness and user’s experience. Therefore, we introduce several multimedia metrics in this paper. Literature in marketing and consumer psychology has already shown that the contextual relevance between the content of hosting webpage and the ads makes a large difference in their clickability [7], and it also has a leading effect on user’s online experience [25]. There exists rich research on measuring the contextual relevance using textual information. Neto et al. [35] investigated ten strategies for a vector space model and evaluated their effectiveness individually. Although the vector space model is widely used, a more advanced matching strategy is required due to the vocabulary impedance problem among ads and webpages. Broder et al. [5] proposed an approach for contextual ad matching based on a combination of semantic and syntactic features. Li et al. [21] further measured the contextual relevance by introducing category relevance, keyword relevance, and style relevance. Category relevance measures whether the ad and webpage belong to the same category; keyword relevance measures the relevance between the keywords extracted from ads and webpage; and style relevance measures the style consistency between the ads and webpage. Zhang et al. [44] proposed an advertising keywords recommendation system for short-text webpages with the help of Wikipedia. Other than textual information, images and videos are more popular on the Internet. Chen et al. [13] allocated ads for images through image recognition techniques. To achieve better contextual relevance, Mei et al. [26, 27] and Guo et al. [16] proposed a multimodal approach which considered both textual and visual relevance for video and image advertising.

Although display banner ads appear on nearly every webpage, their effectiveness remains debatable. Through a series of eye-tracking experiments, recent research found that users tend to avoid ads in web search and surfing [4] and they intentionally avoid looking at such ads even when they are designed to be attention-grabbing [37]. This is also known as *ad overlook*. Moreover, Owen et al. [31] explored the relationship between ads location and the degree of blindness – the phenomenon of website users actively

Table 1: Notations.

Notation	Description
Z, \bar{Z}	Set of auctions
N_z	Set of advertisers in auction z , $z \in Z$
K	Set of metric variables for Stage II re-ranking, $ K = 6$
$b_{i,z}$	Bid price of advertiser i in auction z , $i \in N_z$, $z \in Z$
$b_{(i),z}$	i th highest bid in auction z , $i \in N_z$, $z \in Z$
$y_{i,j,z}$	Probability that advertiser i is allocated to slot j in auction z
$p_{i,z}(b_z)$	Payment of advertiser i in auction z , $b_z = (b_{1,z}, \dots, b_{n_z,z})$
ω_k	Weight of metric variable k in the re-ranking, $k \in K$
$x_{k,*,z}$	Value of metric variable k for the selected advertiser in auction z by our proposed model
$x_{k,\neg,z}$	Value of metric variable k for the selected advertiser in auction z by the SP auction (the ground truth)
$\xi_{k,Z}$	Average change of metric variable k of auctions in set Z
θ_k	Threshold for the change of metric variable k
$rs_{i,z}$	Rank score of advertiser i in auction z
$rs_{*,z}$	Rank score of the selected advertiser in auction z by our proposed model
$rs_{\neg,z}$	Rank score of the selected advertiser in auction z by the SP auction (the ground truth)

ignore web banner ads – and found that users tend to ignore ads located on the bottom and right area. Intuitively, any ad that fails to capture an user’s attention will be ineffective in delivering information. It has been recognized that there are five basic stages before a user accepts a message: attention, comprehension, yielding, retention and action [1]. Motivated by the above observation, in addition to the contextual relevance, we further introduce two other factors to increase user engagement towards banner ads: the visual saliency and the ad image memorability. The former increases the probability that an user will notice the displayed ad while the latter improves an user’s brand perception after she sees the ad image.

3 THE TWO-STAGE FRAMEWORK

The proposed framework consists of two stages. The first stage is based on the second-price (SP) auction model. The second stage re-ranks ads based on the weighted linear combination of benefits of multiple stakeholders, which ensures the selected ad will optimize the trade-offs among them. For the reader’s convenience, the key notations are summarized in Table 1.

3.1 Stage I: Ad Auction

In RTB, the SP auction model has been widely used. Consider auction $z \in Z$, a single impression is bid by n_z advertisers, their cost-per-impression¹ bids are denoted by a bid vector $b_z = (b_{1,z}, \dots, b_{n_z,z})$, and their true values are denoted by a value vector $v_z = (v_{1,z}, \dots, v_{n_z,z})$. In the simplest scenario, the advertiser with the highest bid will be selected and her ad will be displayed to the user. When the ad is displayed, the advertiser then pays the publisher for this page view. In order to re-rank ads in the second stage, we need to know what would be the possible payments from other advertisers. To achieve this, $n_z - 1$ pseudo ad slots are created so that the payments from other advertisers can be estimated.

¹In display advertising industry, the cost-per-impression is usually quoted in terms of cost-per-mille (CPM) which represents the price of 1000 impressions.

We denote the auction outcome by $\{y_{i,j,z}, p_{i,z}\}_{i \in N_z, j \in N_z, z \in Z}$, where $y_{i,j,z}$ is the probability that advertiser i is allocated to slot j in auction z , and $p_{i,z}$ is her payment. Note that slots are ranked so that slot $j \geq 2$ is the created pseudo slot. The auction outcome satisfies the following conditions:

$$0 \leq y_{i,j,z} \leq 1, \quad (1)$$

$$\sum_{i \in N_z} y_{i,j,z} \leq 1, \quad (2)$$

$$\sum_{j \in N_z} y_{i,j,z} \leq 1. \quad (3)$$

Let $b_{(i),z}$ be the i th highest bid in b_z , then

$$y_{i,j,z}(b_z) = \mathbb{I}_{\{b_{i,z} = b_{(j),z}\}}, \quad (4)$$

where $\mathbb{I}_{\{\cdot\}}$ is the indicator function. Then advertiser i ’s payment is

$$p_{i,z}(b_z) = \sum_{j \in N_z} b_{(j+1),z} y_{i,j,z}(b_z), \quad (5)$$

where $b_{(n_z+1),z}$ is equal to the reserve price, whose value can be set to be any value between 0 and $b_{(n_z),z}$.

3.2 Stage II: Optimal Re-Ranking

All advertisers are re-ranked in the second stage. The rank score of advertiser i , $i \in N_z$, is defined as follows:

$$rs_{i,z} = \sum_{k \in K} \omega_k x_{k,i,z}, \quad (6)$$

where $x_{k,i,z}$ is the input value of metric variable k , ω_k is its weight, and $k \in K = \{1, \dots, 6\}$. Below we discuss what are these variables, how to obtain their input values and their weights in the re-ranking.

3.2.1 Specifications of Metric Variables. Table 2 specifies our metric variables in the re-ranking, including the publisher’s revenue, the advertiser’s utility, the ad memorability, the ad CTR, the contextual relevance, and the ad saliency. The input values of these metric variables in the second stage are denoted by $x_{k,i,z}, \forall k \in K$, respectively. They are the normalized values lying between 0 and 1 by using the min-max method for a set of auctions.

Revenue is always the key issue for the publisher. In rank score $rs_{i,z}$, it can be measured by advertiser i ’s (normalized) payment, whose value can be obtained from ad auction in Stage I.

For advertiser i , her benefit can be measured at short-term and long-term levels. The short-term benefit is her utility, which is defined as the difference between her value and payment. It shows the advertiser’s cost saving. In the long run, the ad’s memorability is an important metric, particularly, for branding purpose [28]. It shows how likely the user will remember the advertiser’s ad a few weeks or months later. Here we employ the MemNet model [17] to predict visual memorability of ad image. It is a convolutional neural network (CNN) trained in an annotated image memorability dataset, where the input is an ad image and the output is a single real value which indicates the memorability of this image. The higher the value is, the more memorable is the image.

The user’s benefit can be measured by her page view experience, including the ad CTR, the contextual relevance and the ad saliency [20]. The ad CTR is defined as the number of clicks on the advertiser’s ad divided by the number of displays, whose value is

Table 2: Specifications of metric variables for Stage II.

Variable	Computation			Input for Stage II
	Method	Input	Output	
Publisher's revenue	Stage I	Bid	Payment	$x_{1,i,z} = \text{normalized } p_{i,z}$
Advertiser's utility	Stage I	Bid, payment	Utility	$x_{2,i,z} = \text{normalized } (v_{i,z} - p_{i,z})$
Ad memorability	MemNet [17]	Ad image	Memorability score	$x_{3,i,z} = \text{normalized score}$
Ad CTR	Given by data	CTR	CTR	$x_{4,i,z} = \text{normalized CTR}$
Contextual relevance	TakeLab [38]	Title, keywords, description	Semantic similarity score	$x_{5,i,z} = \text{normalized score}$
Ad saliency	MBS [43]	Web page snapshot, ad image	Saliency map and ratio	$x_{6,i,z} = \text{normalized ratio}$

usually given by data or can be estimated from historical advertising records. In essence, the ad CTR is an ad quality metric – a high CTR means that the advertiser's ad is attractive and also more relevant to the user's needs. The contextual relevance measures if the ad content is more or less relevant to its hosting webpage content. If an ad is more relevant, it will be less intrusiveness [27]. In the paper, we use the TakeLab system [38] to measure the semantic similarity of textual contents between the ad and the webpage. TakeLab uses a support vector regression model with multiple features measuring the word-overlap similarity and syntax similarity. The input is two textual sentences and the output is the semantic similarity score. The score ranges from 0 to 5, where 0 indicates irrelevant and 5 indicates totally relevant. We construct the textual information of the webpage and the ad by using the webpage title, keywords and description, etc. An advantage of TakeLab is that it has an outstanding performance in measuring the similarity between short text snippets, while most other algorithms focus on large documents. The ad saliency metric measures whether the ad image can be easily spotted within its hosting webpage. The more salient ads tend to draw more user's attention [12]. In the paper, we use the minimum barrier salient (MBS) object detection method [43] to calculate the saliency of the ad image. The input is image screenshots of ad and webpage, and the output is the corresponding saliency map. For each pair of webpage and ad candidate, we embed the ad into the webpage and use the MBS method to compute the saliency map and the saliency score is the mean value of each pixel within the ad area.

3.2.2 Determination of the Optimal Weights. The weights of variables in the rank score determine the optimal trade-offs. As described earlier, the publisher needs to sacrifice a certain amount of revenue in the short term in order to increase the benefits of other stakeholders. Let θ_1 be the maximum pre-determined loss rate of revenue and let $\theta_k, \forall k \in K \setminus \{1\}$, be the minimum increase rate of other variables. Given the maximum loss and the minimum increase targets, we can obtain the optimal weights from the training set \bar{Z} .

Our algorithm can be expressed as follows:

$$\max_{\omega_1, \dots, \omega_6} \sum_{z \in \bar{Z}} rs_{*,z}, \quad (7)$$

$$\text{subject to } 0 \leq \omega_k \leq 1, \forall k \in K, \quad (8)$$

$$\sum_{k \in K} \omega_k = 1, \quad (9)$$

$$|\xi_{1,\bar{Z}}| \leq |\theta_1|, \theta_1 \leq 0, \quad (10)$$

$$\xi_{k,\bar{Z}} \geq \theta_k, \theta_k \geq 0, \forall k \in K \setminus \{1\}, \quad (11)$$

where $rs_{*,z}$ is the rank score of the selected advertiser in auction z by our proposed model, $\xi_{k,\bar{Z}}$ is the mean of changes of variable k between our proposed model and the ground truth, defined by

$$\xi_{k,\bar{Z}} = \frac{\sum_{z \in \bar{Z}} (x_{k,*,z} - x_{k,\neg,z})}{\sum_{z \in \bar{Z}} x_{k,\neg,z}}, \quad (12)$$

where $x_{k,*,z}$ is the input value of metric variable k for the selected advertiser in auction z by our proposed model, $x_{k,\neg,z}$ is the input value of metric variable k for the selected advertiser in auction z in the ground truth.

The optimal weights maximize the sum of rank scores of the select advertisers from all auctions in the training set. Eqs. (8)-(9) ensure each variable has an impact in the re-ranking but its impact has an upper bound. Eqs. (10)-(11) further specify the lower bounds of trade-offs where the maximum decrease of the publisher's revenue and the minimum increases for other variables.

3.3 Discussion

Display advertising is an interplay among different stakeholders. A vigorous and healthy advertising eco-system should consider the benefits of all stakeholders. Compared with the existing advertising systems that select ads mainly from the publisher's interest, our proposed framework select ads based on the optimized trade-offs among stakeholders. The incorporated multimedia metrics are computed by the state-of-the-art multimedia techniques, including the MemNet model for the ad image memorability [17], the TakeLab method for the semantic similarity [37], and the MBS method for the visual saliency [43]. Moreover, our framework can be extended with more metrics, such as ad interestingness and image aesthetic.

It should be also noted that the proposed two-stage framework does not ensure advertisers to be truth-telling in Stage I. If an advertiser is truth-telling, her utility can be obtained by computing the difference between her bid and payment. If she is not truth-telling, her value can be learnt from historical ad auctions in Stage I. Value estimation models were discussed by Athey and Nekipelov [2], and Pin and Key [33], respectively. Similarly, we can assume each advertiser is optimal – who maximizes her expected utility in the bidding – and then estimate her value from the training data, and the data should be obtained from the two-stage framework rather than the existing RTB campaigns. On the other hand, we think the proposed framework will not significantly affect an advertiser's truthfulness. Firstly, the three multimedia metric variables and the ad CTR are exogenous. They do not directly affect an advertiser's bidding behaviour. Secondly, advertisers' payments are computed by creating pseudo slots. If an advertiser is not the winning advertiser in the ground truth but she is selected after re-ranking, her

Table 3: Multimedia datasets.

Dataset	AOL	YouTube
Ad network	AOL ONE	Google AdSense
From	06 Sept 2016	06 Sept 2016
To	09 Sept 2016	09 Sept 2016
Location	Singapore	Singapore
# of unique webpages	5,243	7,173
# of unique ads	96	539

Table 4: Ad auction datasets.

Dataset	SSP	Microsoft
Ad type	Display	Search
Ad auction	SP (RTB)	GSP
Market	UK	US
From	08 Jan 2013	26 Dec 2011
To	14 Feb 2013	03 Mar 2012
# of ad slots	31	4,376
# of user tags	NA	NA
# of publishers	NA	1
# of advertisers	374	NA
# of auctions	6,646,643	35,550
# of bids	33,043,127	NA
Bid quote	GBP/CPM	GBP/CPC
Bids of each auction	✓	NA
Winning bid	✓	✓
Winning payment	✓	✓
Estimated CTR	NA	✓

Table 5: Summary of CTRs in the Microsoft dataset.

Position	Minimum of CTR	Maximum of CTR	Mean of CTR	Std. of CTR
1	0.00%	333.30%	13.97%	28.49%
2	0.00%	200.00%	7.97%	15.30%
3	0.00%	750.00%	5.66%	12.47%
4	0.00%	90.90%	4.53%	9.68%
5	0.00%	76.90%	1.53%	3.58%
6	0.00%	333.30%	1.11%	3.62%
7	0.00%	1000.00%	0.68%	5.73%
8	0.00%	108.10%	0.51%	1.65%


actual payment in the two-stage framework is equal to the payment from the SP auction without other advertisers who have higher bids than her. Since advertisers' bids are not disclosed, compared to the existing RTB, the winning advertiser has no further information to support her nontruth-telling biddings.

4 EXPERIMENTS

This section describes our datasets, experimental settings, variable analysis, sensitivity analysis, and overall results.

4.1 Datasets

Four distinct datasets are used in our experiments, including two multimedia datasets and two ad auction datasets. The two multimedia datasets contain the data collected from AOL and YouTube over the period from 06 September to 09 September in 2016. AOL is an article sharing website while YouTube is a video sharing website.

**Figure 2: Unique ads in the multimedia datasets.**

The two platforms use different ad networks: AOL uses AOL Advertising; YouTube uses Google AdSense. Both datasets were collected in Singapore. As we focus on the single slot ad displaying, each webpage in our datasets contains only one ad slot. We start collecting the data from a particular seed URL that contains multiple categories of content.² In this way, the diversity of webpages and banner ads can be ensured. In collecting the data, the web browser is set in the privacy mode, which disables browsing history, web cache and data storage in cookies so that the collected banner ads are not affected by the previous page views. In each dataset, we extract the ads from their webpages to create a set of banner ads and a set of webpages with blank ad slots. For each webpage, the collected data includes title, keywords, description, whole webpage snapshot, ad image. We also crawl title, keywords and description from the ad landing page (i.e., the delivered webpage if an ad is clicked by the visiting user). Note that we do not consider animation ads in our banner ads. When collecting text information about the webpage for YouTube, we also crawl video tags. All the collected text data can be viewed as a summary about the webpage, which can be used to measure the contextual relevance with ad candidates in experiments. Table 3 describes our multimedia datasets. We collect 5,243 unique webpages from AOL and 7,173 unique webpages from YouTube, and the number of their corresponding unique banner ads are 96 and 539, respectively. Fig. 2 shows the growth trends of unique ads on both platforms. Although there are a large number of impressions in the ad network, only a few ads are displayed. We find that some ads re-appear from time to time. The repetitive display strategy reinforces users' memory for branding but it also is a source of intrusiveness into users' online experience [6].

The two ad auction datasets contain campaign information for display advertising and sponsored search. Table 4 provides a brief summary of both datasets. For display advertising, the dataset reports 33,043,127 bids in 6,646,643 auctions over 31 ad slots from a medium-size supply-side platform (SSP) in the UK over the period from 08 January 2013 to 14 February 2013. This SSP dataset has also been used in several recent online advertising studies [8, 11, 42] and is used as the major dataset in our experiments. In the other ad

² AOL: <http://www.aol.com>

YouTube: <https://www.youtube.com/channels>



Figure 3: Relationship among variables in the ground truth data (before normalization).

auction dataset, the estimated campaign results of 547 keywords were crawled from Microsoft adCentre (now Bing Ads) over the period from 26 Dec 2011 to 03 Mar 2012. In adCentre, the targeted keywords, budget, and other settings like matching types were submitted to the system. It then returned a list of estimated statistics, including the estimated cost-per-click (CPC), clicks, and impressions based on 8 positions, which are ranked from top to bottom in the mainline paid listing and sidebar paid listing of search engine result pages. The Microsoft dataset provides information of CTRs, which is important in our proposed model but is not included in the SSP dataset. Table 5 summarizes CTRs across keywords in the Microsoft dataset. Since we focus on the single slot ad displaying, we use the CTR estimated from the first ad position from the mainline paid listing to simulate the CTR in our experiments.

4.2 Experimental Settings

As described in Section 4.1, each of our datasets only provides partial information: the SSP dataset provides advertisers' bids for each auction which can also be used for computing the revenue and the utility; the Microsoft dataset provides CTR distributions

on different ad positions; the AOL and YouTube datasets contain screenshots of ads and webpages and their textual descriptions, which can be used to calculate the multimedia metrics such as ad image memorability, the contextual relevance, and the visual saliency. To simulate a real world display advertising environment, we randomly sample and match records and metrics to connect our auction datasets with the multimedia datasets so that each ad will have both bidding campaign records and multimedia metrics. Given a webpage from the multimedia datasets, we use the following five steps to reproduce its corresponding auction:

Step 1: Sampling an auction in RTB

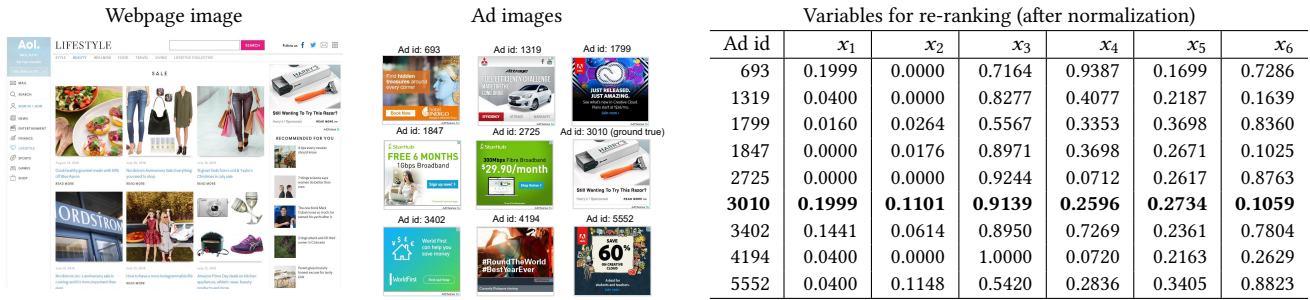
We randomly sample an auction from the SSP dataset. The auction contains a set of advertisers and their bids.

Step 2: Creating the ground truth data

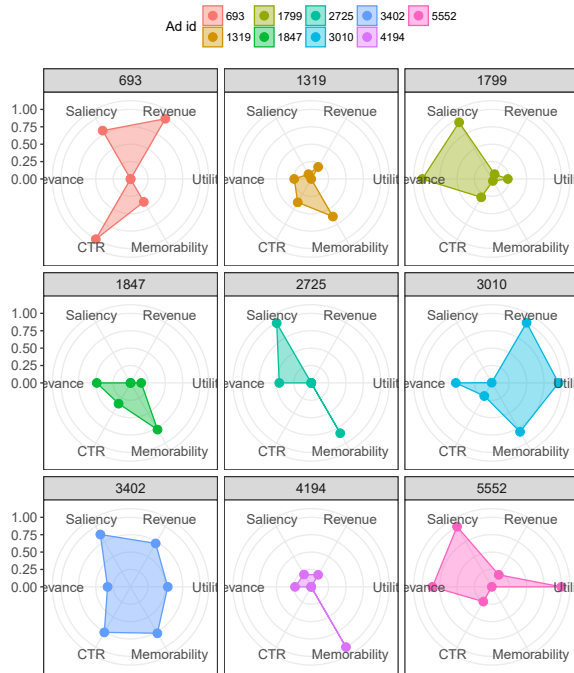
Since the given webpage has an ad displayed when it is web crawled. This original ad is considered as the ground truth.

Step 3: Sampling candidate ads

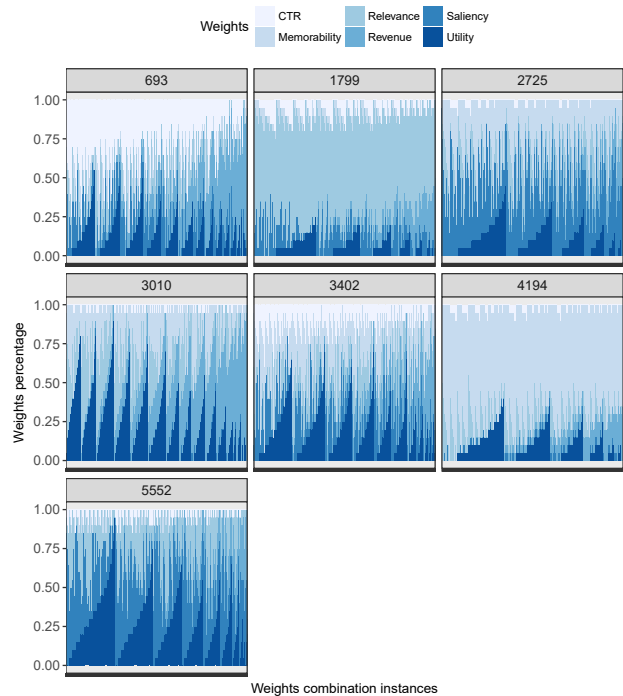
Let n denote the number of advertisers in the sampled auction in Step 1. We then randomly sample $n - 1$ ads from the same



(a)



(b)



(c)

Figure 4: Empirical example of sensitivity analysis of re-ranking for an AOL page view: (a) the raw data in the auction where the ground truth has been highlighted in bold; (b) the comparison of the ground truth ad with other candidates; (c) sensitive analysis of weights in this auction.

multimedia dataset of the webpage. To ensure the uniqueness of each ad in the auction, we exclude the original ad when we sample $n - 1$ ads from the multimedia dataset.

Step 4: Matching ads and bids

We match the n bids with n ads. Since the original ad is displayed, we consider it belongs to the winning advertiser in the auction and match it with the highest bid. The rest $n - 1$ ads and $n - 1$ bids are then randomly matched.

Step 5: Simulating CTRs

For each ad, we generate its CTR from a uniform distribution in the range of 0 and $13.97\% + 1.96 \times 28.49\%$. This is based on the mean and std. of CTRs in the first position in the Microsoft dataset, as also shown in Table 5.

The above five steps allow us to have enough information to compute the mentioned six metric variables for each ad in each auction. We then re-rank ads based on the trade-offs preference, which is expressed in terms of weights. In experiments, we perform the 10-fold cross validation method to obtain the optimal weights as well as to provide the performance analysis.

4.3 Variable Analysis

The values of metric variables in the ground truth are examined – those ads which have been recommended to users under the SP auction model in RTB. For each impression, it contains the publisher’s revenue, the advertiser’s utility, the ad CTR, and the multimedia variables. Fig. 3 shows the descriptive statistics. The density plots

along the diagonal show the distribution of each variable individually; the scatter plots lying in the lower triangular region show how much one variable is affected by another; and their linear correlations are given in the corresponding areas in the upper triangular part. The scatter plots lying on the bottom row shows the values of each variable under the two platforms and the right column show the variations and outliers of each variable. Note that relevance score ranges from 0 to 5, where 5 indicates most relevant and 0 indicates irrelevant, and the scores of memorability and saliency range from 0 to 1, respectively. Several findings are worth mentioning. Firstly, many displayed ads in the ground truth data are irrelevant to the contents of their hosting webpages, which will result in less user engagement. Secondly, most of displayed ads are not salient, thus can be easily overlooked. These two findings may explain the phenomenon of ad overlook in the existing advertising systems [37]. We also find that pairwise variables are independent, which allows using linear combination to calculate the rank score for each ad in the auction without considering the correlation matrix.

4.4 Sensitivity Analysis

Fig. 4 presents an empirical example of analysing trade-offs among stakeholders by investigating all advertisers' ads and possible weights of metric variables. Fig. 4(a) shows the screenshot of the webpage, where ad 3010 (i.e., the ground truth) is displayed on the side bar. Since the sampled RTB auction has 9 advertisers, we further sample 8 candidate ads from the AOL dataset. They have almost the same width and length as the ground truth ad. However, they have different contents, color themes, and descriptions. Fig. 4(a) table exhibits the computed values of variables for re-ranking, where the ground truth has been highlighted in bold. It has the highest value in revenue; and relatively high values in utility, memorability; but small values in CTR, relevance and saliency. Obviously, the ground truth is biased towards the publisher's interest.

Radar charts shown in Fig. 4(b) provide a clearer comparison of the ground truth ad with other ad candidates. Each radar chart has multiple axes and each variable is shown as a point on the axis. A point closer to the center on an axis indicates a lower value and vice versa. In radar charts, all values of variables are further normalized into the range between 0 and 1. If all the variables' values of an ad are less than another ad, we call it being *strictly dominated* because this ad will have no possibility to be selected in the second stage re-ranking. For example, ad 1319 is strictly dominated and its area in the radar chart can be fully covered by ad 3402.

Fig. 4(c) shows all possible combinations of the weights in the re-ranking. For each subplot, x-axis represents all the weights instances that make the ad win the auction, and y-axis represents the values of weights in each instance based on stack bars. Note that the subplots of ad 1319 and ad 1847 are empty, indicating that no matter what weights we choose, these two ads will never be selected in the auction. An ad can be selected if its high valued variables have large weights. Of course, the weights represent the preference of trade-offs and different numbers of weights combination instances of ads also tell that they do not have equal chances to be selected in the re-ranking. Although ad 1847 is not strictly dominated in radar charts, it is not been selected in all weights combinations. This is because most of its variables' values are relatively small and




Figure 5: The effect of θ_1 on the sum of total rank scores of the selected advertisers in auctions.

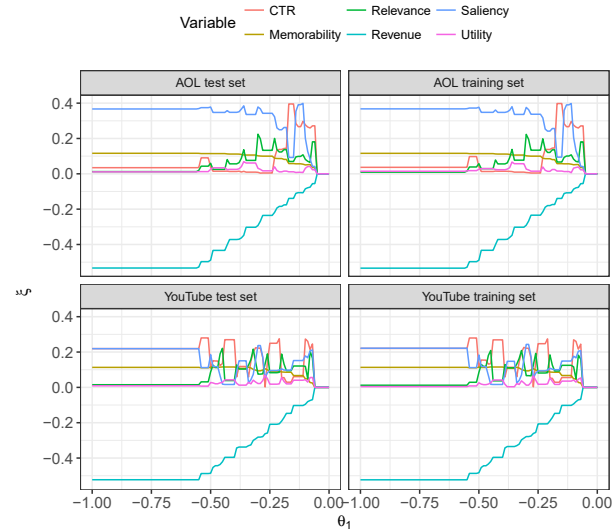


Figure 6: The effect of θ_1 on the changes of variables.

its highest variable value is still less than that of ad 3010. Therefore, even when the trade-offs preference is biased towards the ad memorability, the rank score of ad 1847 is still less than ad 3010. We simply call it to be *weakly dominated*.

4.5 Overall Results

We validate the proposed framework by conducting the 10-fold cross validation on both AOL and YouTube datasets. The optimal weights are estimated from the training set; they are then used for re-ranking ads in the auctions in the test set. For simplicity but without lose of generality, the threshold value θ_k is set to be zero for any $k \in K \setminus \{1\}$ so that the effects of θ_1 can be examined explicitly. As described earlier, θ_1 is the maximum decrease of the publisher's revenue, which indicates how much the publisher would like to sacrifice in order to increase the benefits of the other two stakeholders. Therefore, θ_1 lies in the interval between -1 and 0.

Table 6: Summary of optimal trade-offs among stakeholders in the AOL dataset where $\theta_1 = -0.06$. Note that $\omega_i, i = 1, \dots, 6$ represent the optimal weights obtained from training set, and $\xi_i, i = 1, \dots, 6$ represent the changes of corresponding variables.

Fold	Optimal weight						Training set						Test set					
	ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
1	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	2.3%	3.4%	26.7%	17.4%	4.4%	-5.8%	1.0%	4.0%	30.1%	23.7%	1.3%
2	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	2.5%	3.5%	26.7%	17.6%	4.1%	-5.7%	-0.7%	3.9%	30.6%	22.0%	3.7%
3	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	1.8%	3.6%	27.3%	18.2%	3.7%	-5.7%	6.3%	3.0%	25.3%	16.7%	7.6%
4	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	1.9%	3.5%	27.5%	17.8%	4.1%	-5.1%	4.8%	3.3%	23.4%	19.7%	4.3%
5	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	2.5%	3.4%	27.3%	18.0%	4.1%	-5.5%	0.1%	4.1%	25.4%	18.5%	3.8%
6	0.45	0.30	0.05	0.05	0.15	0.00	-5.6%	2.4%	3.6%	27.2%	18.2%	4.1%	-5.0%	0.9%	2.9%	26.3%	16.2%	4.2%
7	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	2.2%	3.5%	26.8%	18.2%	4.3%	-5.9%	2.5%	3.4%	29.3%	16.7%	2.4%
8	0.45	0.30	0.05	0.05	0.15	0.00	-5.6%	2.1%	3.6%	27.4%	18.4%	4.4%	-4.9%	3.4%	2.6%	24.2%	14.6%	1.5%
9	0.45	0.30	0.05	0.05	0.15	0.00	-5.5%	2.2%	3.5%	27.2%	18.3%	4.0%	-5.3%	2.7%	3.8%	25.7%	15.6%	4.7%
10	0.45	0.30	0.05	0.05	0.15	0.00	-5.4%	2.3%	3.5%	26.6%	18.1%	3.7%	-6.1%	1.3%	4.0%	30.8%	17.0%	8.0%
Mean	-	-	-	-	-	-	-5.5%	2.2%	3.5%	27.1%	18.0%	4.1%	-5.5%	2.2%	3.5%	27.1%	18.1%	4.1%
Std.	-	-	-	-	-	-	0.000	0.002	0.000	0.002	0.003	0.002	0.003	0.020	0.004	0.026	0.027	0.021

Table 7: Summary of optimal trade-offs among stakeholders in the YouTube dataset where $\theta_1 = -0.07$. Note that $\omega_i, i = 1, \dots, 6$ represent the optimal weights obtained from training set, and $\xi_i, i = 1, \dots, 6$ represent the changes of corresponding variables.

Fold	Optimal weight						Training set						Test set					
	ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6	ξ_1	ξ_2	ξ_3	ξ_4	ξ_5	ξ_6
1	0.45	0.25	0.10	0.00	0.15	0.05	-6.2%	0.3%	4.9%	0.0%	23.5%	17.9%	-6.2%	0.5%	4.7%	-2.7%	30.5%	18.2%
2	0.45	0.30	0.05	0.05	0.10	0.05	-6.9%	5.7%	2.2%	27.2%	14.7%	20.9%	-6.8%	7.2%	1.3%	27.4%	18.9%	21.8%
3	0.45	0.30	0.05	0.05	0.10	0.05	-6.9%	5.8%	2.1%	27.4%	14.6%	21.0%	-6.8%	6.8%	2.6%	25.7%	19.2%	21.1%
4	0.45	0.30	0.05	0.05	0.10	0.05	-6.8%	5.4%	2.1%	27.0%	15.4%	21.1%	-7.2%	9.5%	2.4%	29.1%	12.5%	20.1%
5	0.45	0.30	0.05	0.05	0.10	0.05	-6.8%	5.9%	2.1%	27.4%	15.4%	20.9%	-7.2%	5.9%	2.3%	24.9%	12.1%	22.5%
6	0.45	0.30	0.05	0.05	0.10	0.05	-6.9%	6.2%	2.2%	27.3%	15.4%	20.5%	-6.8%	3.4%	2.0%	25.7%	12.4%	25.8%
7	0.45	0.30	0.05	0.05	0.10	0.05	-6.8%	6.0%	2.1%	26.8%	15.1%	20.8%	-7.2%	4.2%	2.4%	30.7%	15.1%	23.1%
8	0.45	0.30	0.05	0.05	0.10	0.05	-6.8%	5.9%	2.1%	27.2%	15.2%	21.1%	-7.2%	5.4%	2.9%	26.6%	14.4%	20.8%
9	0.45	0.30	0.05	0.05	0.10	0.05	-6.9%	5.8%	2.2%	27.1%	14.9%	21.1%	-6.9%	6.3%	1.8%	27.8%	16.5%	20.3%
10	0.45	0.30	0.05	0.05	0.10	0.05	-6.9%	5.8%	2.2%	26.9%	15.3%	21.4%	-6.9%	6.5%	1.8%	30.3%	12.9%	17.7%
Mean	-	-	-	-	-	-	-6.8%	5.3%	2.4%	24.4%	16.0%	20.7%	-6.9%	5.6%	2.4%	24.5%	16.5%	21.1%
Std.	-	-	-	-	-	-	0.001	0.016	0.008	0.081	0.025	0.009	0.002	0.023	0.008	0.092	0.053	0.022

The smaller the value of θ_1 , the more is the increase in the benefits of other stakeholders.

Fig. 5 presents the effect of θ_1 on the sum of rank scores of the selected advertisers in auctions. When θ_1 decreases, the sum of rank scores increases. This monotone decreasing pattern is because that the solution space of weights increases if θ_1 decreases (see Section 3.2.2). Also, in the 10-fold cross validation, the size of the training set is 9 times of the test set, the solution space of weights and the sum of rank scores in the training set are roughly 9 times of the test set. The effects of θ_1 on the changes of metric variables are examined in Fig. 6. In both AOL and YouTube datasets, the changes of variables in the training data follow the same patterns as those in the test data. If θ_1 is close to 0, the SP auction (the ground truth) is the optimal solution. This means that our model suggests that the publisher shouldn't sacrifice any revenue. If the publisher further decreases θ_1 , the optimal weights change. There will be some growth in the variables of other stakeholders, and of course, the publisher's revenue starts to decrease. If θ_1 is small, the publisher's revenue will decrease significantly and she may find it unacceptable. As shown in Fig. 6, if θ_1 is less than -0.3, it reduces the publisher's revenue by almost 30%. There is no monotone increasing pattern in other variables. Analyzing the changes in variables will help the publisher to decide a proper θ_1 . Tables 6-7 further present

two special cases: $\theta_1 = -0.06$ for the AOL dataset; and $\theta_1 = -0.07$ for the YouTube dataset. Both cases show the results when the decrease in the publisher's revenue is close to the pre-specified threshold value θ_1 . For example, there is about 5.5% decrease in the publisher's revenue in the AOL training set and other variables enjoy certain increases ranging from 2.2% to 27.1%.

5 CONCLUSION

This paper discusses a two-stage framework that optimizes trade-offs among stakeholders in display advertising. Different to many related studies which only focus on the publisher's revenue, we further consider the benefits of the advertiser and the user by incorporating multimedia metrics. The trade-offs optimization is based on a linear combination of all metric variables and their weights are learnt from data. Variable and sensitivity analysis are conducted to explore the solution space. Our experimental results validate the proposed framework and show that it is able to increase the benefits of other two stakeholders with just a slight decrease in the publisher's revenue. In the long run, better engagement of advertisers and users will increase the demand of advertising and supply of webpage visits, which can then boost the publisher's revenue. Future research can potentially discuss a framework which considers a generalised scenario where a webpage has multiple ad

slots being separately auctioned off in RTB. The long-term effects of trade-offs optimization can also be further investigated and it would be very interesting if the framework can be tested over a live setting.

6 ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its International Research Centre in Singapore Funding Initiative.

REFERENCES

- [1] L. Alwitt and A. Mitchell. *Psychological processes and advertising effects*. Lawrence Erlbaum, 1985.
- [2] S. Athey and D. Nekipelov. A structural model of sponsored search advertising auctions. In *6th Ad Auctions Workshop*, 2010.
- [3] Y. Bachrach, S. Ceppi, I. Kash, P. Key, and D. Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 75–92, 2014.
- [4] A. Broder, P. Ciccolo, M. Fontoura, E. Gabrilovich, V. Josifovski, and L. Riedel. Search advertising using web relevance feedback. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pages 1013–1022, 2008.
- [5] A. Broder, M. Fontoura, V. Josifovski, and L. Riedel. A semantic approach to contextual advertising. In *Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 559–566, 2007.
- [6] D. Campbell and R. Wright. Shut-up I don't care: understanding the role of relevance and interactivity on customer attitudes toward repetitive online advertising. *Journal of Electronic Commerce Research*, 9(1):62, 2008.
- [7] P. Chatterjee, D. Hoffman, and T. Novak. Modeling the clickstream: implications for web-based advertising efforts. *Marketing Science*, 22(4):520–541, 2003.
- [8] B. Chen. Risk-aware dynamic reserve prices of programmatic guarantee in display advertising. In *Proceedings of the 16th IEEE International Conference on Data Mining Workshops*, pages 511–518, 2016.
- [9] B. Chen and J. Wang. A lattice framework for pricing display advertisement options with the stochastic volatility underlying model. *Electronic Commerce Research and Applications*, 14:465–479, 2015.
- [10] B. Chen, J. Wang, I. Cox, and M. Kankanhalli. Multi-keyword multi-click advertisement option contracts for sponsored search. *ACM Transactions on Intelligent Systems and Technology*, 7(5), 2015.
- [11] B. Chen, S. Yuan, and J. Wang. A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the 8th International Workshop on Data Mining for Online Advertising*, pages 1–9, 2014.
- [12] X. Chen, T. V. Nguyen, and M. Kankanhalli. Salad: A multimodal approach for contextual video advertising. In *2015 IEEE International Symposium on Multimedia*, pages 211–216, 2015.
- [13] Y. Chen, O. Jin, G. Xue, J. Chen, and Q. Yang. Visual contextual advertising: bringing textual advertisements to images. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 1314–1320, 2010.
- [14] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [15] Google. The arrival of real-time bidding and what it means for media buyers. In *White Paper*, 2011.
- [16] J. Guo, T. Mei, F. Liu, and X.-S. Hua. AdOn: an intelligent overlay video advertising system. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 628–629, 2009.
- [17] A. Khosla, A. Raju, A. Torralba, and A. Oliva. Understanding and predicting image memorability at a large scale. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2390–2398, 2015.
- [18] S. Lahaie and P. McAfee. Efficient ranking in sponsored search. In *Proceedings of the 7th International Conference on Internet and Network Economics*, pages 254–265, 2011.
- [19] S. Lahaie and D. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 50–60, 2007.
- [20] J. Lehmann, M. Lalmas, E. Yom-Tov, and G. Dupret. Models of user engagement. In *Proceedings of the 20th International Conference on User Modeling, Adaptation, and Personalization*, pages 164–175, 2012.
- [21] L. Li, T. Mei, X. Niu, and C. Ngo. PageSense: style-wise web page advertising. In *Proceedings of the 19th International Conference on World Wide Web*, pages 1273–1276, 2010.
- [22] W. Liao, K. Chen, and W. Hsu. AdImage: video advertising by image matching and ad scheduling optimization. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 767–768, 2008.
- [23] A. Likhodedov and T. Sandholm. Auction mechanism for optimally trading off revenue and efficiency. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 212–213, 2003.
- [24] B. Lucier, R. Leme, and E. Tardos. On revenue in the generalized second price auction. In *Proceedings of the 21st International Conference on World Wide Web*, 2012.
- [25] S. McCoy, A. Everard, P. Polak, and D. Galletta. The effects of online advertising. *Communications of the ACM*, 50(3):84–88, 2007.
- [26] T. Mei, X. Hua, L. Yang, and S. Li. VideoSense: towards effective online video advertising. In *Proceedings of the 15th ACM International Conference on Multimedia*, pages 1075–1084, 2007.
- [27] T. Mei, L. Li, X. Hua, and S. Li. ImageSense: towards contextual image advertising. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 8(1):6, 2012.
- [28] M. Neumeier. *The Brand Gap: How to Bridge the Distance Between Business Strategy and Design*. New Riders, 2005.
- [29] OpenX. Programmatic + premium: current practices and future trends. Technical report, White Paper, 2013.
- [30] M. Ostrovsky and M. Schwarz. Reserve prices in Internet advertising auctions: a field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 59–60, 2011.
- [31] J. Owens, B. Chaparro, and E. Palmer. Text advertising blindness: the new banner blindness? *Journal of Usability Studies*, 6(3):172–197, 2011.
- [32] D. Parkes. *Algorithmic Game Theory*, chapter Online Mechanisms. Cambridge University Press, 2007.
- [33] F. Pin and P. Key. Stochastic variability in sponsored search auctions: observations and models. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 61–70, 2011.
- [34] F. Radlinski, A. Broder, P. Ciccolo, E. Gabrilovich, V. Josifovski, and L. Riedel. Optimizing relevance and revenue in ad search: a query substitution approach. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 403–410, 2008.
- [35] B. Ribeiro-Neto, M. Cristo, P. Golgher, and E. Silva de Moura. Impedance coupling in content-targeted advertising. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 496–503, 2005.
- [36] B. Roberts, D. Gunawardena, I. Kash, and P. Key. Ranking and tradeoffs in sponsored search auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 751–766, 2013.
- [37] P. Sajjacholapunt and L. Ball. The influence of banner advertisements on attention and memory: human faces with averted gaze can enhance advertising effectiveness. *Frontiers in Psychology*, 5, 2014.
- [38] F. Šarić, G. Glavaš, M. Karan, J. Šnajder, and B. Bašić. Takelab: systems for measuring semantic text similarity. In *Proceedings of the 1st Joint Conference on Lexical and Computational Semantics*, pages 441–448, 2012.
- [39] D. Thompson and K. Leyton-Brown. Revenue optimization in the generalized second-price auction. In *Proceedings of the 14th ACM conference on Electronic Commerce*, pages 837–852, 2013.
- [40] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.
- [41] J. Wang and B. Chen. Selling futures online advertising slots via option contracts. In *Proceedings of the 21st International Conference on World Wide Web*, pages 627–628, 2012.
- [42] S. Yuan, J. Wang, B. Chen, P. Mason, and S. Seljan. An empirical study of reserve price optimisation in real-time bidding. In *Proceedings of the 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1897–1906, 2014.
- [43] J. Zhang, S. Sclaroff, Z. Lin, X. Shen, B. Price, and R. Mech. Minimum barrier salient object detection at 80 fps. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1404–1412, 2015.
- [44] W. Zhang, D. Wang, G. Xue, and H. Zha. Advertising keywords recommendation for short-text web pages using Wikipedia. *ACM Transactions on Intelligent Systems and Technology*, 3(2):36, 2012.