

Joint Optimization of Multiple Performance Metrics in Online Video Advertising

Sahin Cem Geyik *, Sergey Faleev *, Jianqiang Shen, Sean O'Donnell, Santanu Kolay
Turn Inc.
901 Marshall Street, Redwood City, CA, USA
{sgeyik, sfaleev, jshen, sodonnell, skolay}@turn.com

ABSTRACT

The field of online advertising, in essence, deals with the problem of presenting ads to online users in the most appropriate contexts to achieve a multitude of advertiser goals. A vast amount of work in online advertising has been focused on optimizing banner display advertising campaigns where the main goal lies in direct response metrics, often as clicks or conversions. In this paper, we explore the newly popularized space of online video advertising, where brand recognition is the key focus. We propose a framework based on a feedback mechanism where we optimize multiple video specific performance indicators while making sure the delivery constraints (budget and user reach) of advertisers are satisfied. While our main focus is on improving metrics such as engagement (amount of view time), and viewability (whether a campaign is within eyesight of a user), we also discuss the possibilities of expanding to other metrics. We demonstrate the benefit of our framework via empirical results in multiple real-world advertising campaigns. To the best of our knowledge, this is the first paper that deals with the unique challenges arising from the nature of online video advertising.

CCS Concepts

•Information systems → Computational advertising;
Expert systems;

Keywords

Online Advertising; Video Advertising; Performance Optimization

1. INTRODUCTION

Online advertising aims to match advertisers with the most beneficial online users through the most appropriate

contexts (websites, mobile applications etc.) in order to maximize their campaign specific goals. These campaign specific goals are often one of two kinds: brand recognition metrics, i.e. any metric which serves to indicate whether or not an audience's opinion of an advertiser's product has improved, or whether or not the size of the audience that knows about an advertiser (or a specific product) has increased; and direct response metrics, which can be directly measured in terms of monetary value, such as a click or a conversion (product purchase, subscription to a newsletter, etc.) [11].

The online advertising ecosystem is generally accepted to be comprised of components which collectively enable advertisers to reach their aforementioned goals: *Demand-side Platforms* (DSPs) strive to find opportunities on behalf of advertisers to show an ad (impression) to a user beneficial to a specific advertiser; *Supply-side Platforms* (SSPs) are publishers (websites or mobile application providers) or intermediary platforms (which publishers connect to) that supply advertising opportunities for multiple competing DSPs (often in a second-price auction) to capture and show an advertisement on; *Data Management Platforms* (DMPs) deal with processing first-party and third-party data for advertisers in order to enable business and analytical inferences for more effective advertising.

By the help of the above three components, the online advertising domain has experienced tremendous growth in the past few years, both in terms of the quality of ads and the amount of Internet traffic allocated to online advertising [19]. In terms of the media type used for advertising, banner display ads have been the main industry standard for the past two decades [4], which are graphic ads designed to achieve user attention, and whose components can also be calibrated for further optimization [9]. In accordance with this trend, a vast amount of literature has been focused towards optimizing direct response metrics which are clicks [6] or conversions [11], since it is not trivial to estimate brand recognition metrics for banner ads. This is especially true for evaluating whether the users' opinions of an advertiser improved or not, which is even more complex.

In this paper, our efforts are mainly focused on optimizing video advertising campaigns where the ads shown to users are in the form of videos, which can have a variety of durations. The key difference for this type of advertising (compared to display ads) is that brand recognition metrics are used much more commonly compared to direct response metrics. Our focus is on two metrics whose detail is given in the next section: (i) *Engagement*, which depends on the

* The authors contributed to this work equally.

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.

KDD '16, August 13-17, 2016, San Francisco, CA, USA

© 2016 ACM. ISBN 978-1-4503-4232-2/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2939672.2939724>




Figure 1: Spend distribution over metric priorities for banner display advertising.

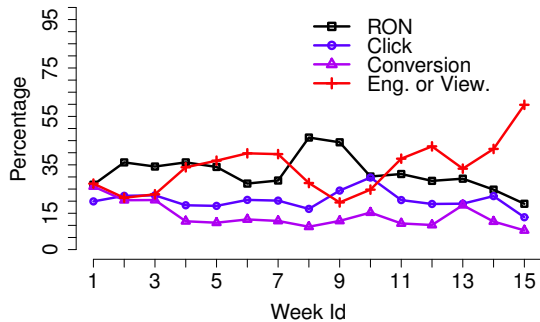


Figure 2: Spend distribution over metric priorities for video advertising.

length of time a video ad is watched, and (ii) *Viewability*, which depends on the location of the video ad inside a webpage or an application. We also propose a feedback-based mechanism that both optimizes the aforementioned two metrics, and satisfies the advertisers’ delivery constraints. To the best of our knowledge this is the first work that deals with the specific challenges of online video advertising.

The main contributions of this work include:

- A detailed examination of challenges that arise from online video advertising,
- A multi-criteria optimization framework specific to video advertising which aims to also satisfy advertisers’ delivery constraints, and,
- Real-world experiments and results that demonstrate the advantages of our approach.

2. BACKGROUND AND CHALLENGES OF VIDEO ADVERTISING

As aforementioned, we explore the details of video advertising in this work, which is inherently different compared to banner display advertising. First of all, the advertisers that follow the path of display advertising* often utilize click or conversion optimization campaigns. While the advertisers also wish to receive brand recognition through run-of-network campaigns (RON) where no optimization is

* Within this paper, from time to time, we will use the term *display advertising*, for short, instead of *banner display advertising*.

implied with fixed bid values, aiming to explore new audiences via relaxed user targeting constraints, they have no way to measure whether the users were indeed interested in their product, except for the direct response metrics. Please see Figure 1 for the distribution of campaign spending for display advertising on Turn’s platform for a 15-week period within 2015 over different metric prioritizations which supports our claim. In online video advertising however, the advertisers are mainly concerned with brand recognition. This makes it necessary to use metrics which are more closely tied to evaluating the brand recognition performance of a campaign. In this paper we are mainly focusing on three such metrics (the first one specific to video advertising):

- **Engagement** metric measures the percentage of the total duration of a video ad that has been actively watched by the user. There can be multiple types of engagement events (e.g. 25%, 50%, 75%). In the context of our current work, we take engagement to be a binary event where 1 represents that the user watched the whole (100%) of the video ad, and 0 otherwise. Please note that, intuitively, engagement metric measures users’ opinion of an advertiser, since it indicates the steady interest (as measured by the amount of time ad was watched) of the user towards the ad.
- **Viewability** indicates the goodness of the ad’s location, hence its *viewability* by the online user. Formally, the *view* event of an ad impression is 1 if the video ad is within eye-sight of the user during some time interval (industry standard is 2 seconds with at least 50% of ad area is within eye-sight). This metric suggests size of the audience knowing about an advertiser or product, since the knowledge of a user on an advertiser has not formed if the user has not *viewed* the ad.
- **User Reach** indicates how many users have watched the ads. Formally, since it is common for advertisers to put a frequency cap on the number of impressions delivered to one single user, *User Reach* can be calculated as the number of delivered impressions.

Figure 2 shows the distribution of spending for the video campaigns within Turn that signals importance on different metrics for the same 15-week period as in Figure 1. It can be seen that there is a clear upwards trend of advertisers utilizing viewability and engagement optimization tactics. Furthermore, we can also observe that the amount of budget spent for conversion campaigns holds a much smaller portion of the total video spend compared to display campaigns.

Lack of interest in direct response metric for video advertising has both advantages and disadvantages. A significant advantage comes from the fact that several problems arising for campaigns that value conversions are no longer of concern. First of all, one of the greatest problems related to conversion prediction/tracking is delayed feedback [2], since it takes a significant amount of time for a user to act (in terms of *conversion*, i.e. product purchase etc.) after seeing an ad, or clicking it. Secondly, tracking conversions requires a methodical way to attribute the conversion to the campaigns that show the ad to the user [17, 5], which is not trivial. This means that when we have a campaign that prioritizes conversions as its performance metric, there will be times when it will not even be able to register those conversions (because they have been attributed to another campaign) for accurate performance estimation.

The main disadvantage of having campaigns that prioritize brand recognition (not direct response) metrics is the ambiguity in assigning a monetary value to the evaluated metrics. While it is trickier for an advertiser to give a monetary value for a click, it is often trivial to see how much monetary value an advertiser should assign to a conversion. For example, the amount of money spent by the user when s/he purchases a product, or the profit made by the advertiser for that conversion, can yield a very intuitive monetary value which can be assigned to a certain conversion. This way, we can look at return-on-investment (ROI) for a specific campaign, which is the ratio of monetary value received (via clicks or conversions) by the campaign to the total amount of money spent for impressions by the campaign. Therefore, in some cases, we do not even have to have a great reach (number of unique users which are shown an impression) or delivery (the amount of money spent for advertising via impressions), as long as we can get a sufficient number of direct responses. This however is not the case for brand recognition metrics. There is no trivial way to assign monetary value to an engagement or view event by a user. Hence, the problem we need to solve for online video advertising includes both optimization of brand recognition metrics, as well as satisfying user reach and delivery criteria set by the advertiser. Please note that satisfying the delivery constraints of an advertiser campaign by itself is a difficult problem which is severely affected by the characteristics of the audience that the advertiser wants to reach [7, 18]. This difficulty gets further aggravated in video advertising campaigns where we aim to improve brand recognition metrics as well.

As mentioned in the previous section, coming up with a system that optimizes brand recognition metrics as well as delivery is our focus. After the literature survey in the next section, we will first delve into the formal definition of the problem we are trying to solve. We then will present our proposed methodology which aims to perform a constant trade-off between the delivery of a campaign and engagement as well as view events received by that campaign.

3. PREVIOUS WORK

Based on the domain-specific goals and metrics of success for online video advertising, we have introduced the challenges that are significantly different from those of the current mainstream approach, which is display advertising. As expected, display advertising has also been much more extensively studied in literature. To the best of our knowledge, our current work is the first to address the specific challenges arising in online video advertising. However, considering that our work also employs a joint optimization over various metrics, we will detail prior research done and approaches taken to jointly optimize performance and/or delivery metrics in online advertising. Please note that there are very few works detailing optimization efforts in online *video* advertising, and these in general exist in the substantially different context of online video advertising infrastructure. One such work deals with jointly optimizing the allocation of video rates and inserted advertisement duration in a wireless broadcast network to maximize profit [8].

Online Video Advertising: There have been a variety of proposed systems to implement effective forms of online video advertising, focused on insertion scheduling, video context similarity, and user relevancy issues. Inserting advertisements into a video stream is an especially complicated

problem, as video advertisements are likely to degrade user experience if done clumsily, reducing the effectiveness of a campaign and potentially driving away users. Intuitively, advertisements are less disruptive when they are contextually relevant, with one proposed system [12] using advertiser-provided images (relevant to their own brand/message) to be compared against sampled video frames to determine the best candidate advertisements to show in a given video.

Another previous effort [14] focused on detecting sets of non-intrusive advertisement insertion points within a video, in addition to using contextual and visual-aural relevance to select the best advertisements to place. The authors demonstrated that this approach improved user viewing experience. A similar work [20] chose to introduce the ability to bookmark or delay interaction with an advertisement, in order to avoid video viewing interruption partially or entirely. Finally, in [22], the authors focus on *user's attention relevancy*, where assuming a video ad triggered by a user can be considered relevant to a user's attention, they propose a framework that generates semantic concepts from video broadcasts, which are subsequently used to rank video advertisements by relevancy concepts to find the most relevant video advertisements to ones triggered by a user previously. This paper presented findings which suggested that such an approach was regarded as less intrusive and more satisfying.

Compared to previous work in online advertising, our current effort differs in multiple ways. First of all, the listed literature is mostly involved with systems for scheduling and presenting video content, and have access to the actual video content to parse and analyze ahead of time; this is not applicable to a system like ours, which participates in the online advertising auction process in real-time under strict constraints. Additionally, most previous work is evaluated based on user studies and subjective measures, without much of a feedback loop or evaluation metric to evolve their systems' algorithms or models over time. As a DSP, we are essentially solving a different set of problems in online video advertising, which are necessarily complicated by the evaluation metrics that we are striving to satisfy.

Joint Optimization in Online Advertising: In our current proposal, we are not claiming to introduce new methods of joint optimization, but rather, demonstrating a novel application for joint optimization of desired metrics in online video advertising and its empirical effectiveness. To the best of our knowledge, our proposed use of joint optimization over multiple brand recognition metrics for online video advertising is a novel effort, whereas joint optimization techniques have seen historical use in more traditional areas of online advertising. Perhaps the closest such use to our approach is a framework for the joint optimization of guaranteed delivery (often a brand advertiser's goal) and pay-per-click (a performance-centric goal) auctions in sponsored search, with the objective of increasing search engine revenue [16]. This effort demonstrated the effectiveness of joint optimization compared to many baseline strategies. Additional efforts in sponsored search included jointly optimizing budget allocation and bid price setting, while presenting that it outperformed both bid price optimization and campaign budget optimization done separately [24].

A multi-objective programming approach [1] to optimizing click-through rate and downstream engagement metrics (e.g. a third-party reported time spent on a page), with constraints on how much click volume can be lost, is analo-

gous to our approach, except for direct-response metrics and advertisers. Finally, previous research at the intersection of joint optimization applications and online video advertising seems to exist in a complete different facet of the online advertising industry (in this case, the actual provisioners of video traffic, i.e. the wireless video broadcast network itself) [8], and is not applicable to our challenges.

4. RESEARCH PROBLEMS

Optimization Goals of Video Advertising. Real-time Bidding enables advertisers to target their ads to specific users and cherry-pick only the impressions they deem most valuable to them. The core of our system is a bidding algorithm that computes an appropriate bid for an ad to be submitted to a webpage or mobile app that is being presented to an online user. There is a plethora of efforts in the literature addressing the problem of *bid estimation*. For traditional display advertising, bid optimization can be modeled as a linear programming problem where the objective is to maximize the expected revenue and the constraint is the budget of an advertiser [3]. The bid or the cost per impression is then modeled as:

$$\text{bid} = P(\text{outcome} \mid \text{user, page, ad}) \times \text{Goal}, \quad (1)$$

where the outcome is usually a click or conversion (i.e. direct responses, which we mentioned previously to be the main focus of display advertising) and the *Goal* is defined by the advertiser and indicates the monetary value of such an outcome. *User* and *page* dimensions are provided by the incoming bid request. The majority of research is focused on estimating the probability of a click or conversion. Researchers typically employ logistic regression and its variants or frequency-based features to estimate this probability [11, 15, 6, 13].

Applying the above approach directly to video advertising is not feasible. First of all, as we explained in sections 1 and 2, video advertising campaigns mainly deal with brand-recognition metrics as a priority, which are harder to assign a monetary value. Furthermore, instead of having one single goal, video advertisers typically have multiple performance metrics that can be challenging to be optimized simultaneously. Such metrics can be calculated in multiple ways, depending on the advertiser's preferences. Popular ones include *engagement rate* (ER), *viewability rate* (VR), *user reach* (UR), *cost per click* (CPC), *cost per action* (CPA), and *return-on-investment* (ROI). Assuming we get T bid requests $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$, we can calculate the performance of a bid model f_A for advertiser \mathcal{A} as follows:

$$\text{ER}(f_A, X) = \frac{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{complete}(\mathcal{A}, \mathbf{x})}{|\text{win}(f_A, X)|},$$

$$\text{VR}(f_A, X) = \frac{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{viewable}(\mathcal{A}, \mathbf{x})}{|\text{win}(f_A, X)|},$$

$$\text{UR}(f_A, X) = |\text{win}(f_A, X)|,$$

$$\text{CPC}(f_A, X) = \frac{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{cost}(f_A, \mathbf{x})}{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{click}(\mathcal{A}, \mathbf{x})},$$

$$\text{CPA}(f_A, X) = \frac{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{cost}(f_A, \mathbf{x})}{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{conv}(\mathcal{A}, \mathbf{x})},$$

$$\text{ROI}(f_A, X) = \frac{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{value}(\mathcal{A}, \mathbf{x})}{\sum_{\mathbf{x} \in \text{win}(f_A, X)} \text{cost}(f_A, \mathbf{x})}. \quad (2)$$

Above,

- $\text{win}(f_A, X)$ denotes the set of bid requests that are won by model f_A and have been shown some impressions of \mathcal{A} , and $|\text{win}(f_A, X)|$ is the number of impressions shown to users by advertiser \mathcal{A} ,
- $\text{complete}(\mathcal{A}, \mathbf{x})$ indicates whether the impression of \mathcal{A} on \mathbf{x} was watched by the user for its whole duration, and $\text{viewable}(\mathcal{A}, \mathbf{x})$ indicates whether the impression of \mathcal{A} on \mathbf{x} is within eye-sight of the user reported by the third-party,
- $\text{click}(\mathcal{A}, \mathbf{x})$ is a binary indicator on whether the impression on \mathbf{x} leads to a click for \mathcal{A} , and $\text{conv}(\mathcal{A}, \mathbf{x})$ indicates whether the impression on \mathbf{x} leads to a conversion for \mathcal{A} ,
- $\text{cost}(f_A, \mathbf{x})$ is the amount of money spent by model f_A to win bid request \mathbf{x} , and $\text{value}(\mathcal{A}, \mathbf{x})$ is the amount of value (by the clicks and conversions performed) provided by winning bid request \mathbf{x} to advertiser \mathcal{A} .

Optimizing Multiple Goals. Given advertiser \mathcal{A} with budget B and optimization goals $\langle \phi_1, \phi_2, \dots, \phi_L \rangle$ where each element is one of the metrics defined in Equation 2, our ultimate objective can be formulated as in the following:

DEFINITION 1. The video bidding optimization algorithm tries to search for a model f that calculates and submits bids for requests $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ for advertiser \mathcal{A} , so that

$$\begin{aligned} \max_{f: \mathbf{x} \rightarrow \mathbb{R}} & \langle \phi_1(f, X), \phi_2(f, X), \dots, \phi_L(f, X) \rangle \\ \text{s.t.} & f(\mathbf{x}_t) \geq 0 \quad , \forall 1 \leq t \leq T, \\ & \sum_{\mathbf{x} \in \text{win}(f, X)} \text{cost}(f, \mathbf{x}) \leq B, \end{aligned} \quad (3)$$

In other words, given the constraints on budget and inventory, we seek an optimal bidding strategy to maximize multiple returns for our customers.

DEFINITION 2. A bid model f is sound if it satisfies both budget and inventory constraints; a sound model $f^1: \mathbf{x} \rightarrow \mathbb{R}$ dominates another sound model $f^2: \mathbf{x} \rightarrow \mathbb{R}$ iff

$$\phi_l(f^1, X) \geq \phi_l(f^2, X) \quad , \forall 1 \leq l \leq L;$$

A sound model f^* is utopian iff f^* dominates any other sound model f .

In most cases, given the optimization goals $\langle \phi_1, \phi_2, \dots, \phi_L \rangle$, it is infeasible to find a *utopian* bidding model for the advertiser. For example, assume advertiser \mathcal{A} wants to optimize both the click rate and the engagement rate. It is clear that they can be contradictory to each other: to maximize clicks, we want to make the ads as eye-catching as possible; however, once the ads get attention, it is inevitable that some users will skip the video play. On the other hand, the video ads could finish without skipping (100% engagement event) if they are invisible but in such cases the click rate will be hurt.

Since *utopian* bidding models rarely exist for video advertising, searching for appropriate models is not as straightforward as for the conventional optimization problems of

display advertising. In some cases, the advertiser is neutral about the priority of optimization goals and does not require an absolute ranking order. We can then search for the optimal bidding model by formulating a single-objective optimization problem such that the optimal solution to the single-objective optimization problem cannot be dominated by any sound model of the multi-objective optimization problem. This can be achieved by proper scalarization of the objectives:

$$\begin{aligned} \max_{f: \mathbf{x} \rightarrow \mathbb{R}} \quad & w_1 \cdot \phi_1(f, X) + w_2 \cdot \phi_2(f, X) + \dots + w_L \cdot \phi_L(f, X) \\ \text{s.t.} \quad & f(\mathbf{x}_t) \geq 0 \quad , \forall 1 \leq t \leq T , \\ & \sum_{\mathbf{x} \in \text{win}(f, X)} \text{cost}(f, \mathbf{x}) \leq B , \end{aligned} \quad (4)$$

Here, w_l is the advertiser's preferences for different optimization goals. By choosing different parameters for the scalarization, we can produce bidding models with emphases on different goals. However, in many cases the above scalarization approach is not applicable due to two reasons: first, many advertisers have sufficient preference and the optimization goals can be ranked in the order of importance; second, the choice of weight w_l has a significant impact on the generated bidding model and it is not straightforward for advertisers to choose weighting parameters that make the most sense to them.

In this paper, we focus on the case where an advertiser's optimization goals can be ranked in order of importance. Let's assume, without loss of generality, that the performance goals are in an order of importance so that ϕ_1 is the most important and ϕ_L is the least important to the advertiser. Then searching for the optimal bidding models becomes solving a sequence of L single-objective optimization problems of the form

$$\begin{aligned} \max_{f: \mathbf{x} \rightarrow \mathbb{R}} \quad & \phi_l(f, X) \\ \text{s.t.} \quad & \phi_j(f, X) \geq y_j^* \quad , \forall 1 \leq j \leq l-1 , \\ & f(\mathbf{x}_t) \geq 0 \quad , \forall 1 \leq t \leq T , \\ & \sum_{\mathbf{x} \in \text{win}(f, X)} \text{cost}(f, \mathbf{x}) \leq B , \end{aligned} \quad (5)$$

where y_j^* is the optimal performance value of the above problem with $l = j$. In other words, in each iteration we estimate the optimal performance value and use it as the constraint in the next iteration. Let us give an example to make this optimization scheme clearer - suppose we solve the optimization problem for the first objective, and we find f_1 which gives the optimal $\phi_1(f_1, X) = y_1^*$. Then, in the next iteration, we will find the best f_2 to optimize $\phi_2(f_2, X) = y_2^*$, such that f_2 also satisfies $\phi_1(f_2, X) \geq \phi_1(f_1, X) = y_1^*$. We obtain the final bidding model after the L^{th} problem is solved. Please note that this optimization scheme only makes sense if the function f_l we find in each iteration is not unique, otherwise we will not update the solution for the rest of the iterations (no other f will be feasible for any other iterations). Practically, the nonuniqueness of bidding models f_l could be realized as a quasi-random (with respect to optimization goals ϕ_{l+1} , ϕ_{l+2} , etc) choice of the bidding opportunities in which the ad actually participates when the budget could be satisfied by bidding on a fraction of the available to a bidding opportunity.

In this paper, we will discuss how to approximately solve

the above problem and calculate the bidding model for video advertising when an advertiser's optimization goals are strictly ranked in order of importance.

Special Case of Dual Goals. In general, the most important goal of display advertising is to acquire as many clicks/actions as possible. It is different for online video advertising: like traditional TV advertising, its major purpose is branding – utilizing video impressions to create a differentiated name and image of products in order to establish a presence in the consumer's mind, hence attracting and retaining customers. To achieve this, advertisers usually want to maximize user reach (UR) and deliver as many impressions as possible. Besides maximizing user reach, video advertisers also tend to want a secondary performance goal, which can further indicate brand-recognition, such as engagement rate (ER) and viewability rate (VR). Although less common, as aforementioned, some advertisers do value direct response metrics in video advertising, such as cost per click (CPC) and cost per action (CPA). Thus, the most common case for optimizing video advertising has dual goals: one is UR and another is ER/VR/CPC/CPA.

As an emerging media, the winning price of video advertising impressions is usually high and predictable [23]. Given that an advertiser's budget is limited, maximizing impressions is then equivalent to spending all of the available budget. From this it follows that the optimization problem of dual goals given in (5) can be formulated as:

$$\begin{aligned} \max_{f: \mathbf{x} \rightarrow \mathbb{R}} \quad & \phi(f, X) \\ \text{s.t.} \quad & f(\mathbf{x}_t) \geq 0 \quad , \forall 1 \leq t \leq T , \\ & \sum_{\mathbf{x} \in \text{win}(f, X)} \text{cost}(f, \mathbf{x}) = B , \end{aligned} \quad (6)$$

Here, ϕ is the secondary (after UR) metric of importance, which are VR, ER, CPC, or CPA. The bidding model needs to spend the entire budget to ensure maximum audience reach. However, solving the above problem directly is challenging: first of all, the expected value for a bid request (\mathbf{x}_t) is hard to forecast, especially for a newly created ad; second, the huge size of bid request patterns, as well as the real-time nature (as we do not have a global view of the bid requests, and these come in a sequence) of online advertising, make direct search prohibitively expensive. It is necessary to adopt an efficient approximate solution, which we will introduce in the next section. Our methodology relies on a feedback mechanism that constantly changes the manner with which we optimize the secondary goal, based on the current status of the UR satisfaction (since this is our first priority).

5. METHODOLOGY

In order to improve both user reach and other performance metrics (e.g. viewability rate, engagement rate, CPC, or CPA), we adopt an approximation of the sequential goal optimization. Within this approach, advertisers are allowed to set different priorities for different metric goals. The goal with the highest priority is always assigned to the user reach, hence our algorithm first tries to satisfy the user reach goal (in terms of budget to be delivered). Once we estimate that the primary goal can be satisfied, the algorithm tries to improve the second priority goal (e.g. viewability rate), and this process keeps iterating until we find a balance point where the goals are optimized. In each step, the algorithm

tries to adjust the bidding model while keeping the higher priority goal satisfied. Such adjustments end when the metric with lower priority cannot be improved without trading-off the metric with higher priority.

5.1 Framework of Dual Optimization

As discussed above, we first focus on the case of dual goals and later describe case with more than two goals. Our approach to solving the dual optimization problem can be considered as a special case of *reinforcement learning* [21]: the bidding model makes a bid based on the system dynamics, hoping to achieve optimal rewards; based on the received feedback, we might need to constantly adjust the bidding model. It is a continuous evolutionary process in the Markovian environment. Like other reinforcement learning problems, we rely on accurate modeling of the environment to improve the secondary goal of a campaign while optimizing the primary goal. Having such models allows the system to generate virtual experience and predict the consequences of actions before they are taken. Our models of the world consist of two components:

- **Estimation of Impression Need** β , which is a numerical value between 0 and 1, where the larger value indicates more desire for impressions. We take the historical bid traffic and win rate as input to calculate the delivery expectation in real-time for the campaign. Such signals enable the platform to understand how well the acquisition of impressions goes,
- **Estimation of Goal Metric** ϕ , which maps each perceived $\langle \text{BidRequest}, \text{Bid} \rangle$ pair to a single number, a *reward*, indicating the intrinsic desirability of the bid. We apply machine learning algorithms on historical bidding data in order to train hierarchical prediction models [11].

Impression Need β for a campaign indicates its user reach satisfaction, and our system then cherry-picks bid requests according to their expected performance for ϕ . Here, we give an illustrative example before the formal explanation of our algorithm. Suppose that the advertiser sets the primary campaign goal to be user reach (UR), and the secondary goal to be viewability rate (VR). As we noted above, the user reach goal will be satisfied if Impression Need $\beta < 1$ (ideally, $\beta \ll 1$). The viewability rate goal is satisfied by introducing a dynamic filtering threshold for the viewability rate, θ_v . For each bid request, the website and user information is extracted from the bid request parameters, and expected viewability rate ϕ_v is estimated for the qualified campaign utilizing the $\langle \text{ad}, \text{website}, \text{user} \rangle$ combination. The campaign will not submit a bid for the request if $\phi_v < \theta_v$.

The framework for calculating the dynamic filtering threshold θ is given in Algorithm 1, which performs the secondary goal optimization, while satisfying the primary goal of user reach. For every campaign, θ is calculated periodically in real-time together with β . We start with a conservative θ at the beginning of the campaign to give it a boost in spending. After a new β is calculated for a given campaign, the θ increases accordingly if β is in the *comfortable delivery* mode region, and decreases if β is in the *risky delivery* mode region, and does not change otherwise. These regions (the function *classify* in Algorithm 1) are calculated based on the hypothesis checking for upper and lower limits for

Algorithm 1: Optimization of Dual Goals

Input: θ : Current threshold for the secondary goal
Output: θ' : New threshold for the secondary goal

- 1 $\beta \leftarrow$ calculate Impression Need ;
- 2 $\text{status} = \text{classify}(\beta)$;
- 3 $\Delta \leftarrow$ calculate the moving step based on β ;
- 4 **if** *status is Comfortable Delivery* **then**
- 5 $\theta' = \theta + \Delta$;
- 6 **else if** *status is Risky Delivery* **then**
- 7 $\theta' = \theta - \Delta$;
- 8 **else**
- 9 $\theta' = \theta$;
- 10 $\theta' = \min(1, \max(\theta', 0))$;
- 11 **return** θ'

each region. If the user reach or delivery situation is constantly good (β is in the comfortable delivery region), then the secondary goal is improved in a methodical manner by consistently increasing the threshold.

5.2 Estimating Impression Need

We need a proxy variable in real-time to evaluate how well we are satisfying the user reach goals, which is our first priority, of a campaign. Before explaining the design of Impression Need, we would like to first explain the life cycle for a bid request, which ends in an advertisement shown (which is the ‘impression’ event) to the user. When a Real-Time Bid exchange (RTB exchange, which is a part of the SSP component of the online advertising ecosystem) sends a bid request to the DSP, several campaigns are usually qualified (through passing targeting criteria set by the advertisers) for bidding. These campaigns participate in an internal auction performed by the DSP, where they are evaluated based on multiple criteria set by the advertisers, and the advertisement belonging to the campaign with the highest score (for CPA or CPC campaigns this score is proportional to the value returned by Equation 1) is sent (together with the bid price in the bid response) back from the DSP. The bid responses sent to the RTB exchange by several DSPs participate in the RTB Exchange’s auction (which is usually *second price*), where the ad with the highest bid wins the auction and is served to the online user.

Usually advertisers want their campaign budgets to be delivered fully with more or less steady pacing within a day (similar amounts of user reach achieved for different periods of the day) [10]. In such a case, in order to ensure steady pacing of the budget, the campaign only participates in β fraction of internal DSP auctions for which it qualifies (according to the targeting criteria set by the advertiser). This dynamic parameter β [$0 \leq \beta \leq 1$] directly reflects an advertiser’s need for impressions. We heuristically adjust β at time period $t + 1$ based on “ideal” pacing of the remaining daily budget, the amount of user reach (hence budget spent) in the previous time interval, and traffic:

$$\beta_{t+1} \propto \beta_t \frac{B_{t+1}}{S_t} \frac{\text{BidRequest}(t)}{\text{BidRequest}(t+1)} \frac{\text{WinRate}(t)}{\text{WinRate}(t+1)} \quad (7)$$

where B_{t+1} is the budget to be spent at $t + 1$, S_t is the spending at t , BidRequest is the number of bid requests (we calculate $\text{BidRequest}(t+1)$, i.e. number of bid requests

that is predicted to be achieved in time $t+1$, using historical data) and $WinRate$ is the average probability of winning a bid. Impression Need, β , can be applied to ensure smooth budget pacing within a day while finishing the daily budget by the end of the day (this would indeed qualify as satisfying the user reach, UR, as explained in § 4). The case when $\beta \ll 1$ indicates a good delivery situation (i.e. the campaign participates in only a very small fraction of qualified internal DSP auctions, because it is comfortably satisfying its UR constraint), and the case when $\beta \approx 1$ (the campaign participates in all qualified internal auctions) indicates a poor delivery situation, when the campaign has the risk of not satisfying budget constraints in (5) and (6).

5.3 Improving Brand Recognition

Please note that the framework given in Algorithm 1 improves the secondary goal if and only if we have a comfortable delivery pattern for the campaign. When the campaign struggles to attain its delivery or user reach goals, we need a methodical way to increase the bids for the campaign, since higher bid amounts increase a campaign’s chances of winning the RTB exchange’s auctions.

LEMMA 1. *When the budget is not fully spent, higher bid prices will obtain more impressions.*

PROOF. Real-time bidding follows a second price auction mechanism, therefore increasing the bid prices would not affect our cost on those impressions that had been acquired before. On the other hand, given higher bid prices, some bids that we lost before could have been won. \square

This in a way causes the campaign to pay higher than its expected value for each impression; however, it also helps the campaign to satisfy its primary goal, which is user reach. There is clearly a trade-off between the primary and secondary goals of a campaign. In this case, our proposed approach calculates the bid (b) (for the bid request that this campaign qualifies for) as a function of three parameters, ϕ_v , θ_v , and the maximum average bid price specified by the advertiser, b_m :

$$b = f(b_m, \phi_v, \theta_v), \quad (8)$$

where b_m is provided for a campaign when the advertiser does not want to pay more than b_m per served impression on average.

For proprietary reasons and company policy, we are not able to give the exact details of the function f , however, in general, $f(b_m, \phi_v, \theta_v)$ monotonically increases with increasing ϕ_v or b_m . We would like to emphasize here that this function $f(b_m, \phi_v, \theta_v)$ is used for brand-recognition metrics (e.g. viewability and engagements) when there is no trivial way to assign a monetary value for the events seen.

We utilize the threshold θ_v as the parameter that describes current delivery situation, instead of *Impression Need*, β , in f . This is due to the statistical nature of Impression Need which is caused by the internet traffic. We have observed that β could fluctuate in a broad range within a day, while the threshold θ_v incrementally changes with time and therefore has a more smooth behavior that reflects the delivery trends at a larger time scale.

An illustrative example of f is shown in Figure 3. In a comfortable delivery situation ($\theta_v > 0$), the bidding function $f(b_m, \phi_v, \theta_v)$ decreases in the region of small and intermediate ϕ_v , therefore shifting the average in-view rate to

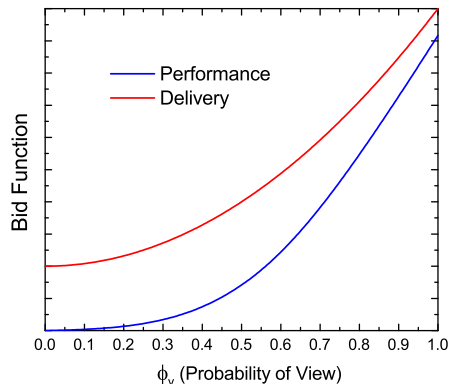


Figure 3: Illustration of bidding functions of dynamic model in performance and delivery limits.

higher values for won impressions. This has a similar effect of cherry-picking bid requests with predicted high viewability. In risky delivery situations ($\theta_v \approx 0$), the function $f(b_m, \phi_v, \theta_v)$ increases in the region of small and intermediate ϕ_v , and eventually, at $\theta_v = 0$ reaches maximum value of b_m for $\phi_v = 0$. Please note that $f(b_m, \phi_v, \theta_v \approx 0)$ still increases for increasing ϕ_v , hence improving the secondary goal of the campaign. Such behavior reflects the fact that the advertiser’s primary goal is user reach or delivery, and if delivery is bad, the bidding function should be as high as possible (restricted only by maximum average bid price), even for bid-requests with a small probability of success. Note that for $\phi_v > 0$, the bidding function could be larger than b_m only if the average price per impression for the given campaign over several days is less than b_m , hence we have some relaxation in bidding.

In the above example, the secondary goal is completely ignored (bid is simply $b = b_m$) if the campaign has extremely tight targeting and therefore the delivery goal is hard to satisfy ($\beta = 1$ and therefore $\theta_v \approx 0$) and also if the average price per impression in last several days is close to the limit of b_m . On the other hand, if delivery is good ($\beta \ll 1$), the secondary goal is fully considered and θ_v approaches the maximum possible value (constrained by delivery staying in a neutral region), hence cherry-picking good bid requests to improve the secondary goal.

5.4 Improving Direct Response

Note that the bid function f described in the previous subsection is for brand recognition metrics (i.e. viewability and engagement). This approach significantly differs from the bid function for conventional CPC or CPA optimization goal types. For these goal types, the advertiser sets the maximum average bid price, b_m , as well as the dollar goal value, *Goal*, which indicates the monetary value each click or action is worth. In such a case, due to the second price nature of RTB exchange auctions, the optimal strategy is to bid as given in Equation 1. We can still use the methodology given in § 5.1 to receive more clicks and conversions per ad shown to online users by a campaign. However, modification of the shape of the bidding function of Equation 1 is suboptimal since it may either reduce the return on investment (ROI) below 100% in the case of bidding higher than optimal, or lead to possible underdelivery (user reach goal is not satisfied) in the case of bidding lower than optimal.

In the case of the viewability or engagement goal types, there is no monetary goal value assigned to a view or an engagement event, and therefore there is no consistent ROI definition. This means that we can utilize different shapes of the bidding function due to what the advertiser may prefer to optimize towards - higher delivery or higher performance. The extreme case of the delivery optimization bidding function is the constant bidding function, $b = b_m$, whereas the extreme case of the performance optimization bidding function is a function that is very small in the region of small and intermediate ϕ_v , and increases only in the region of large probabilities of success, $\phi_v \approx 1$.

As we detailed in § 5.3, based on the current delivery situation (described by the dynamic filtering threshold θ_v), our bidding function dynamically changes in a range - from a constant bidding strategy, i.e. $b = b_m$, in the case of bad delivery, to the secondary goal optimizing bidding strategy in the case of good delivery. As we will demonstrate in the experiments section, this dynamic bidding approach turns out to give a better combination of delivery and performance compared to static bidding functions that are set towards the optimization of either delivery or performance.

5.5 Beyond Dual Goals

Currently we implemented the above goal optimization algorithm with only first and second priority goals. The expansion to larger than two goals makes sense only if the targeting criteria set by the advertiser for a campaign is loose enough compared to its budget (user reach or delivery goal), therefore enough inventory is available for the campaign to enable optimization of more than two goals. While this is an open problem, we will sketch here an idea which may indeed help with optimizing towards multiple goals while the user reach constraint is achieved.

In the case of more than two optimization goals, we believe that the bidding function should be chosen based on the second priority goal. If the second priority goal is viewability or engagement, then the bidding function from Equation 8 is used. If the secondary goal is a direct response metric such as CPC or CPA, then the bidding function of Equation 1 is used with dynamic thresholding for *click-through rate*, θ_c , or *conversion rate*, θ_a . After the second priority goal is improved, i.e. if we have high values of θ_2 , we can utilize Impression Need β to optimize towards the next priority goals by utilizing θ_i where $i > 2$. (Here θ_i is the dynamic threshold for goal i .) At any time period when Impression Need enters a risky delivery region, the thresholds for the lower priority goals should be reduced first until the comfortable delivery region is reached, obeying the goal priorities set by the advertiser.

Practically, the system should estimate some maximum threshold set, $\hat{\theta}_i < 1$, such that achieving $\theta_i \approx \hat{\theta}_i$ will be considered as satisfaction of the goal i . This allows the system to progress to the next optimization goal. Indeed, for dual-goal optimization with set $\hat{\theta}_2$, then $\theta_2 = \hat{\theta}_2$ and $\beta \ll 1$ could be the final equilibrium state for well-spending campaigns, while without θ_2 , θ_2 will increase until reaching the limit of $\theta_2 \approx 1$ while β being still in comfortable delivery zone (but not $\beta \ll 1$ anymore).

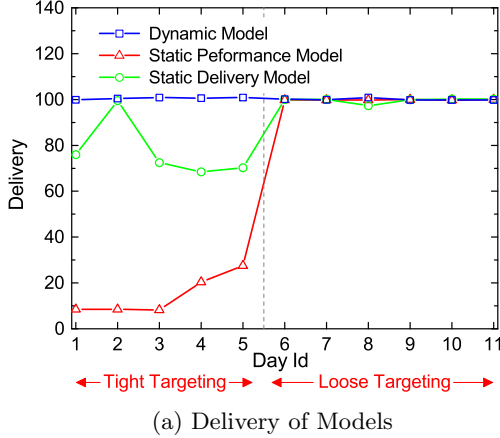
6. EXPERIMENTAL RESULTS

In this section, we will present some preliminary results on real-world campaigns that employ the filtering and bid-

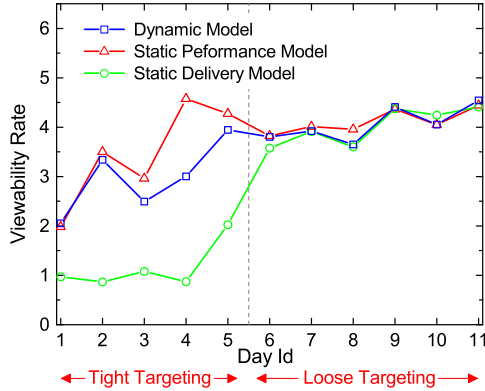
ding methodologies listed in § 5. In Figure 4, we demonstrate the delivery and viewability rate for three models: (i) dynamic model, (ii) static performance model, and (iii) static delivery model. Among these, the *dynamic model* is the sequential goal optimization scheme described in § 5, where the primary goal is delivery and secondary goal is the viewability rate. The bidding function of this model is dynamically adjusted based on the current delivery situation (which is approximated by the proxy variable θ_v , viewability rate filtering threshold). The two static models given in the figure are the viewability optimization models that are similar to the dynamic model in all aspects (i.e. they have the same mechanism for the cherry-picking of bid requests with threshold θ_v as the dynamic model) except that the bidding function is static, hence it does not depend on θ_v (bids do not change with the current delivery situation, i.e. how well we are on track to satisfy the delivery or user reach constraint). The static performance model has a bidding function which is optimized for higher performance (we only bid high when the predicted viewability rate is significantly high), while the static delivery model has a bidding function which is optimized for higher delivery (aims for delivery by bidding higher even for bid requests with low estimated viewability). The bidding function of the static performance model can be given as $f_p(b_m, \phi_v) = f(b_m, \phi_v, \theta_v = 1)$ (from § 5). The bidding function of the static delivery model, $f_d(b_m, \phi_v)$, has larger values than $f_p(b_m, \phi_v)$ in the region of small and intermediate ϕ_v (while still $f_d(b_m, \phi_v = 0) = f_p(b_m, \phi_v = 0) = 0$), but these functions have similar values in the region of large ϕ_v .

The dynamic and the two static models were run for the same campaign and the same daily budget during an 11-day period within 2015 (we replicated all the properties of these campaigns, such as targeting, while only tweaking the bidding strategy). We applied a tight targeting constraint on all the campaigns for the first five days to study the performance of all the listed models in a potentially risky delivery situation. In the last six days, a loose targeting tactic has been applied to the campaigns to see the outcome of the models in a comfortable delivery setting. It can be observed from Figure 4(a) that in the first five days, the delivery of the dynamic model is better than both of the static models*. In the tight targeting situation, the dynamic model fully delivers its daily budget while the static delivery model delivers about 80% of its daily budget, and the static performance model delivers less than 30%. While it seems counter-intuitive to see the dynamic model deliver better than the static *delivery* model, this behavior is expected due to the utilized bidding functions. The dynamic model bids higher (especially in the region of small ϕ_v) than *both* of the static models in a bad delivery situation. Figure 4(b) shows viewability rate for the three models, where we can observe that the dynamic bidding model performs nearly as good as the static performance model. In summary, it is apparent that the dynamic model spends better than both the static models (and, in particular, much better

* Please note that due to the company's privacy policy, we have applied a normalizing factor to all the results presented in this section. The *Delivery* axis in the graphs represents the monetary amount spent on bid requests by each of the campaigns, which is limited by their daily budget.



(a) Delivery of Models




(b) Viewability Rate of Models

Figure 4: Delivery and viewability rate of the dynamic and static models applied to the same campaign with the same daily budget for 11 days. For the first 5 days, tight targeting has been applied on the campaign, and for the last 6 days, loose targeting has been applied.

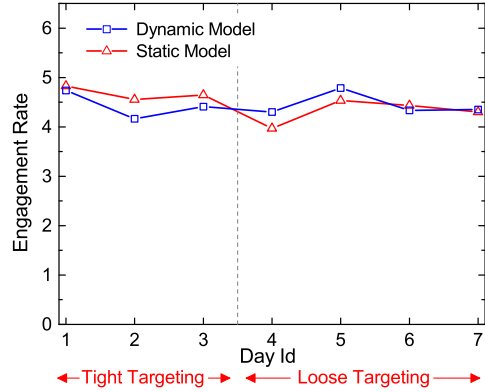
than the static performance model) and provides a viewability rate very close to the static performance model in the tight targeting case.

If we examine the results from the last six days when the loose targeting criteria were applied to the campaigns, we can see from Figure 4(a) that all three campaigns performed in a very similar manner. This is a consequence of the fact that, in a comfortable delivery situation, the filtering threshold θ_v is high and all three models only bid on bid-requests with large ϕ_v . Since the bidding functions of all three models behave similarly in the region of large ϕ_v , the delivery and viewability rates of all three models are similar.

We present the delivery as well as the engagement rates of the dynamic model and the baseline static engagement model (which does not adjust its bid values, but rather only applies filtering through θ_e , i.e. engagement rate filtering threshold) in Figure 5 for a 7-day period within 2015. Again, we replicated the same campaign with the same daily budget, where only the bidding functions differ. The primary goal of the dynamic model is delivery, and the secondary goal is the engagement rate. The bidding function of the dy-



(a) Delivery of Models



(b) Engagement Rate of Models

Figure 5: Delivery and engagement rate of the dynamic and the static models applied to the same campaign with the same daily budget for 7 days. For the first 3 days, tight targeting has been applied on the campaign, and for the last 4 days, loose targeting has been applied.

amic model is similar to the bidding function of the static model for intermediate and large values of ϕ_e (predicted engagement rate for a bid request). In the region of small ϕ_e , the bidding functions differ. In particular, similar to the viewability case, the bidding function of the dynamic model is equal to b_m in the limit of $\theta_e = 0$ and $\phi_e = 0$, while the bidding function of the static model (which is independent of θ_e) is equal to 0 in the limit of $\phi_e = 0$.

For the engagement model evaluation campaigns, we applied three days of tight targeting, and four days of loose targeting, to see the effects of both comfortable and risky delivery situations. Since the targeting criteria applied to the identical campaigns were extremely tight during the first three days, both models could not satisfy their daily budget, yet the dynamic model still delivered twice as much advertising budget due to bidding higher on bid-requests with small ϕ_e . During the last four days, when looser targeting criteria had been applied to the campaigns, both models spent their daily budget and showed similar engagement rates. We can observe from Figure 5(b) that engagement rates stayed pretty much the same during the whole seven days, includ-

ing the tight targeting period. This shows that the engagement rate for the user set reached by the campaigns was quite homogeneous, and further demonstrates that even by increasing the bid values and delivering higher amounts of budget, the dynamic model indeed did not hurt performance in terms of engagement rate.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we focused on the recently popularized domain of online video advertising. We listed the challenges and peculiarities related to video advertising which are inherently different compared to the more traditional banner display advertising approach. We also proposed a multi-objective optimization scheme for online video campaigns, which has a compelling theoretical foundation in its problem definition but is also quite challenging to implement considering the real-time nature of online advertising. We presented a methodology to approximately perform this joint optimization, and demonstrated via real-world experiments that the proposed framework improves delivery and brand recognition metrics, such as viewability and engagement rates for video campaigns. To the best of our knowledge, this is the first work to examine the difficulties that arise due to the nature of online video advertising and present ways to improve advertiser metrics specific to online video campaigns.

As mentioned in the introduction of our methodology, the first future work we would like to propose is extending the number of prioritized goals as opposed to the dual-objective optimization we focused on in this paper. We have presented an initial idea on this work, but nevertheless, this is an interesting direction to follow. Furthermore, we did not perform a deep-dive into the problem of predicting viewability or engagement rates of bid requests for video campaigns. It is our belief that the effective methods to be applied for this purpose could prove to be radically different compared to the previous efforts in literature on click and conversion estimation, and would be an exciting subject to pursue.

Acknowledgments

We thank many talented scientists and engineers at Turn for their help and feedback in this work.

8. REFERENCES

- [1] D. Agarwal, B.-C. Chen, P. Elango, and X. Wang. Click shaping to optimize multiple objectives. In *Proc. ACM KDD*, pages 132–140, 2011.
- [2] O. Chapelle. Modeling delayed feedback in display advertising. In *Proc. KDD*, pages 1097–1105, 2014.
- [3] Y. Chen, P. Berkhin, B. Anderson, and N. R. Devanur. Real-time bidding algorithms for performance-based display ad allocation. In *Proc. ACM KDD*, pages 1307–1315, 2011.
- [4] D. S. Evans. The online advertising industry: Economics, evolution, and privacy. *The Journal of Economic Perspectives*, 23:37–60, 2009.
- [5] S. C. Geyik, A. Saxena, and A. Dasdan. Multi-touch attribution based budget allocation in online advertising. In *Proc. ACM ADKDD*, pages 1–9, 2014.
- [6] X. He et al. Practical lessons from predicting clicks on ads at facebook. In *Proc. ADKDD*, pages 1–9, 2014.
- [7] A. Jalali, S. Kolay, P. Foldes, and A. Dasdan. Scalable audience reach estimation in real-time online advertising. In *Proc. ICDMW*, pages 629–637, 2013.
- [8] W. Ji, Y. Chen, M. Chen, B.-W. Chen, Y. Chen, and S.-Y. Kung. Profit maximization through online advertising scheduling for a wireless video broadcast network. *IEEE TMC*, 2015.
- [9] G. G. Karuga, A. M. Khraban, S. K. Nair, and D. O. Rice. Adpalette: An algorithm for customizing online advertisements on the fly. *Decision Support Systems*, 32:85–106, 2001.
- [10] K.-C. Lee, A. Jalali, and A. Dasdan. Real time bid optimization with smooth budget delivery in online advertising. In *Proc. ACM ADKDD*, pages 1–9, 2013.
- [11] K.-C. Lee, B. Orten, A. Dasdan, and W. Li. Estimating conversion rate in display advertising from past performance data. In *Proc. ACM KDD*, pages 768–776, 2012.
- [12] W.-S. Liao, K.-T. Chen, and W. H. Hsu. Adimage: Video advertising by image matching and ad scheduling optimization. In *Proc. ACM SIGIR*, pages 767–768, 2008.
- [13] H. B. McMahan et al. Ad click prediction: a view from the trenches. In *Proc. ACM KDD*, pages 1222–1230, 2013.
- [14] T. Mei, X.-S. Hua, L. Yang, and S. Li. Videosense: Towards effective online video advertising. In *Proc. ACM MM*, pages 1075–1084, 2007.
- [15] C. Perlich, B. Dalessandro, R. Hook, O. Stitelman, T. Raeder, and F. Provost. Bid optimizing and inventory scoring in targeted online advertising. In *Proc. ACM KDD*, pages 804–812, 2012.
- [16] K. Salomatin, T.-Y. Liu, and Y. Yang. A unified optimization framework for auction and guaranteed delivery in online advertising. In *Proc. ACM CIKM*, pages 2005–2009, 2012.
- [17] X. Shao and L. Li. Data-driven multi-touch attribution models. In *Proc. ACM KDD*, pages 258–264, 2011.
- [18] J. Shen, S. C. Geyik, and A. Dasdan. Effective audience extension in online advertising. In *Proc. ACM KDD*, pages 2099–2108, 2015.
- [19] J. Shen, B. Orten, S. C. Geyik, D. Liu, S. Shariat, F. Bian, and A. Dasdan. From 0.5 million to 2.5 million: Efficiently scaling up real-time bidding. In *Proc. IEEE ICDM*, 2015.
- [20] S. H. Srinivasan, N. Sawant, and S. Wadhwa. vadeo: Video advertising system. In *Proc. ACM MM*, pages 455–456, 2007.
- [21] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 1998.
- [22] J. Wang, Y. Fang, and H. Lu. Online video advertising based on user’s attention relevancy computing. In *Proc. IEEE ICME*, pages 1161–1164, 2008.
- [23] W. C.-H. Wu, M.-Y. Yeh, and M.-S. Chen. Predicting winning price in real time bidding with censored data. In *Proc. of ACM KDD*, pages 1305–1314, 2015.
- [24] W. Zhang, Y. Zhang, B. Gao, Y. Yu, X. Yuan, and T.-Y. Liu. Joint optimization of bid and budget allocation in sponsored search. In *Proc. ACM KDD*, pages 1177–1185, 2012.