



Blacklisting the Blacklist in Digital Advertising

Improving Delivery by Bidding for What You Can Win



AdKDD2017

Yeming Shi, Claudia Perlich, Ori Stitelman

[yshi, claudia, ostitelman@dstillery.com](mailto:yshi,claudia,ostitelman@dstillery.com)

Dstillery, Inc.



Outline

- Introduction
- Bidding Process
- Publisher-Marketer Blacklist Identification Algorithm
- Out-of-time Validation of the Algorithm
- In Vivo Performance Results
- Optimal Amount of Exploration

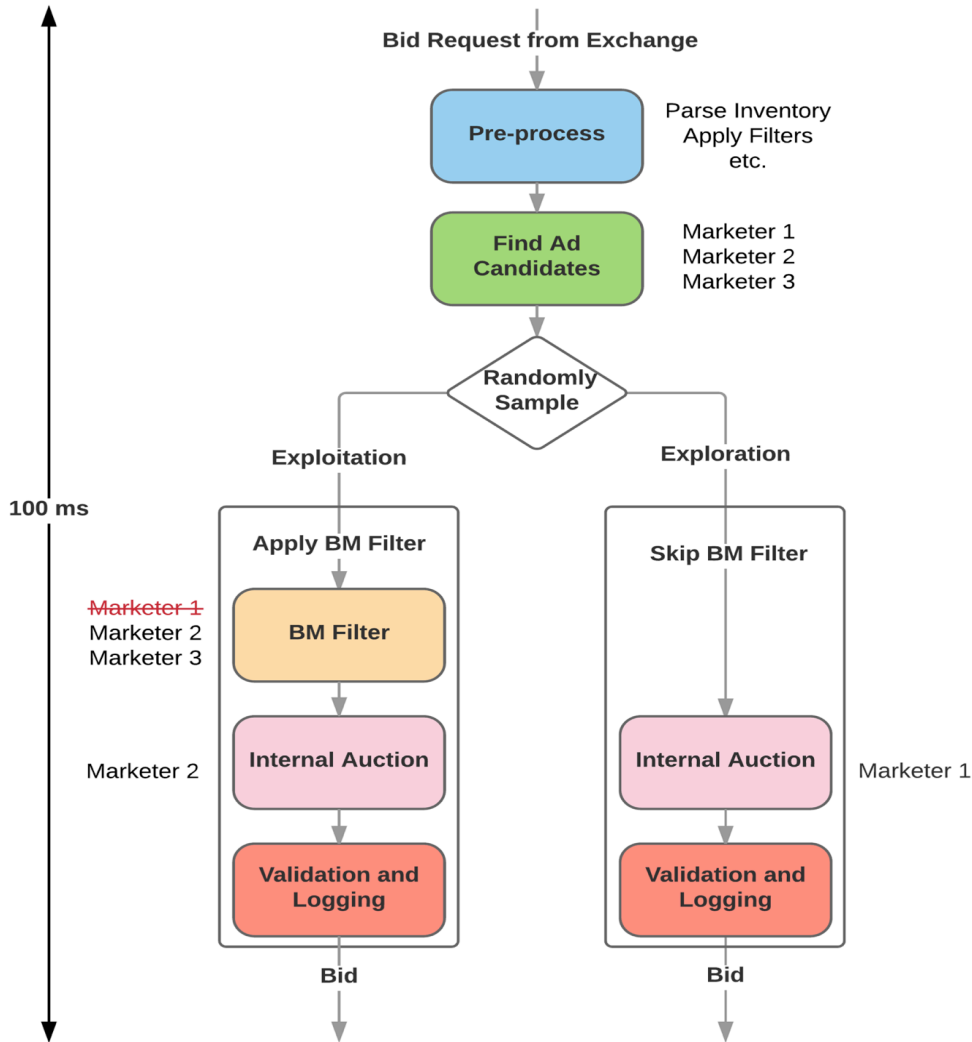
Digital Advertising: \$50 Billion Industry

- Billions of ad slots bought & sold per day on exchanges through real time auctions
- Buyers bid to show ad on publisher (inventory) on behalf of brands or marketers.
- For each received bid request (BRQ), a buyer needs to decide which marketer to bid for.
- The decision is based on the opportunity's values (click-through rate or conversion rate) across multiple ad campaigns.
- These values are estimated using machine learning.

Motivation

- It is commonly believed that the highest bid wins the ad auction, but there are exceptions.
- A publisher sometimes bans a marketer for example to avoid conflict with direct deals.
- Buyers bidding for a banned marketer (BM)
 1. waste computing resources in a low latency environment
 2. lose of an opportunity to show a good ad for a different marketer
- At Dstillery, we've built a system to automatically detect the content of a publisher's blacklist and limit bidding for BMs.

Bidding Process



Marketers 1, 2, and 3 are ad candidates in descending order of bid price.

Left path: Marketer 1 in blacklist (BM list) => bid for marketer 2.

Right path: A small sample of BRQs skip the BM filter (exploration).

BM Identification Algorithm

- Idea: if a marketer's win rate (WR) is much lower than other marketers' WR on the same exchange and publisher => this marketer is a BM
- We also need to account for bid price because WR is sensitive to price, otherwise we could misidentify a marketer as a BM because its bid price is low.
- How? Discretize bids into price buckets and compare actual number of impressions against number of impressions expected from other marketers' WR.

Algorithm

e: exchange i: inventory m, m': marketer

p: price bucket

Bids: # of bids

(Tot)Imps: # of (total)

Impressions

- Step 1: Discretize all bids into price buckets for each exchange / inventory. Compute expected WR of marketer m as the average WR of all marketers:

$$\mathbb{E} [WR_{e,i,m,p}] = \frac{\sum_{m'} Imps_{e,i,m',p}}{\sum_{m'} Bids_{e,i,m',p}} .$$

- Step 2: Compute number of expected impressions for marketer m in each price bucket

$$\mathbb{E} [Imps_{e,i,m,p}] = Bids_{e,i,m,p} \cdot \mathbb{E} [WR_{e,i,m,p}] .$$

- Step 3: Aggregate all price buckets to get expected total impressions and actual total impressions.

$$\mathbb{E} [TotImps_{e,i,m}] = \sum_p \mathbb{E} [Imps_{e,i,m,p}] ,$$

$$TotImps_{e,i,m} = \sum_p Imps_{e,i,m,p}$$

- Step 4: Put Marketer m in BM list if its actual total impressions is much less than expected total impressions

$$TotImps_{e,i,m} < \alpha \mathbb{E} [TotImps_{e,i,m}] - \beta \sqrt{\mathbb{E} [TotImps_{e,i,m}]}$$

Algorithm

e: exchange i: inventory m, m': marketer

Bids: # of bids

p: price bucket

(Tot)Imps: # of

(total) impressions

BM Identification Criterion: actual total impressions much less than expected total impressions

$$TotImps_{e,i,m} < \alpha \mathbb{E}[TotImps_{e,i,m}] - \beta \sqrt{\mathbb{E}[TotImps_{e,i,m}]}$$

Choice of (α, β) : a business not statistical problem.

A practically good choice is (0.6, 1.2) .

Requires enough data (expected total impressions ≥ 7.0) for effective detection.

- Dynamic update: every day, refresh BM list based on past $T = 8$ days of data and apply filter in bidding.

Out-of-time Validation of the Algorithm

Generate the BM list based on past 8 days' data (estimation set).
On the next day (test set), Monitor WRs for (exchange, inventory, marketer) combinations in and outside the BM list.

Evaluation Set	in BM list	Bids	WR	$\mathbb{E}[WR]$
Estimation	False	$1.24e^9$	5.99%	5.31%
	True	$2.31e^8$	0.81%	4.49%
Test	False	$1.22e^8$	6.37%	5.75%
	True	$2.34e^7$	0.94%	4.12%

Validation Results:

- No over-fitting
- Changes in WR are not too fast to render the learning ineffective.

In Vivo Performance Results

- WR lift:
Control group: apply BM filter
Treatment group: no BM filter
WR in control / treatment groups following deployment:

Day	Control Group	Treatment Group	Lift
Day 1	3.37%	4.17%	1.24
Day 2	3.23%	3.96%	1.23

Daily WR lift [1.23 ± 0.03](#) since adoption (4 months)

- Not bidding on combinations in the BM list results in a 20~30% system load reduction in terms of the number of bids.

In Vivo Performance Results

- When a BM is filtered, there is a 48% chance that we find a qualified alternative marketer to bid. This results us delivering 6% more impressions while making 14% fewer bids.
- Finding an alternative marketer is good for alternative's campaign performance(clicks or conversions). Our pacing engine activates just the right amount of the highest-ranked opportunities to meet each campaign's delivery goal. In the absence of BM filter, we lose active opportunities to show good ads to the alternative marketer. Consequently, the pacing engine has to activate a lower ranked opportunity to deliver for the alternative, which hurts alternative's performance.

Exploration-Exploitation Balance

- Dedicate a fraction (c) of bids to exploration.
- Tradeoff between large vs. small control group sampling rate c :
 1. Large $c \Rightarrow$ bid too much for BMs
 2. Small $c \Rightarrow$ bid too little on some blacklisted (exchange, inventory, marketer) combinations \Rightarrow can't tell they're blacklisted later \Rightarrow end up bidding more on them
- An optimal c minimizes the number of bids spent on experimenting with BMs.

Computing Optimal Exploration Sampling Rate for One BM

- Approach: compute the number of bids made per day for one BM combination as a function of c , then minimize w.r.t. c .
- To detect a BM, need at least N bids made for it in the past $T = 8$ days.
- N is related to the BM identification criterion and our global WR. $N = 175$.
- Three classes of BM combinations :
 1. Small BM: never caught in list
 2. Median BM: enters and exits list periodically, forming $(T+1)$ -day cycles
 3. Large BM: always in list

Number of Bids Made Per Day for One BM

Number of bids made per day for a BM that we naturally (in absence of BM filter) bid for n times per day:

$$b_{e,i,m} = \begin{cases} n_{e,i,m} & \text{if } n_{e,i,m} < N/T \\ c \cdot n_{e,i,m} & \text{if } n_{e,i,m} \geq N/(c \cdot T) \\ b_I(n_{e,i,m}) & \text{otherwise} \end{cases} ,$$

where

$$b_I(n_{e,i,m}) = \frac{n_{e,i,m} \cdot (T + 1 - x) + c \cdot n_{e,i,m} \cdot x}{T + 1} ,$$
$$x = 1 + \left\lfloor \frac{T - N/n_{e,i,m}}{1 - c} \right\rfloor .$$

x is the number of days a median size BM remains in BM list over a $(T+1)$ -day cycle.

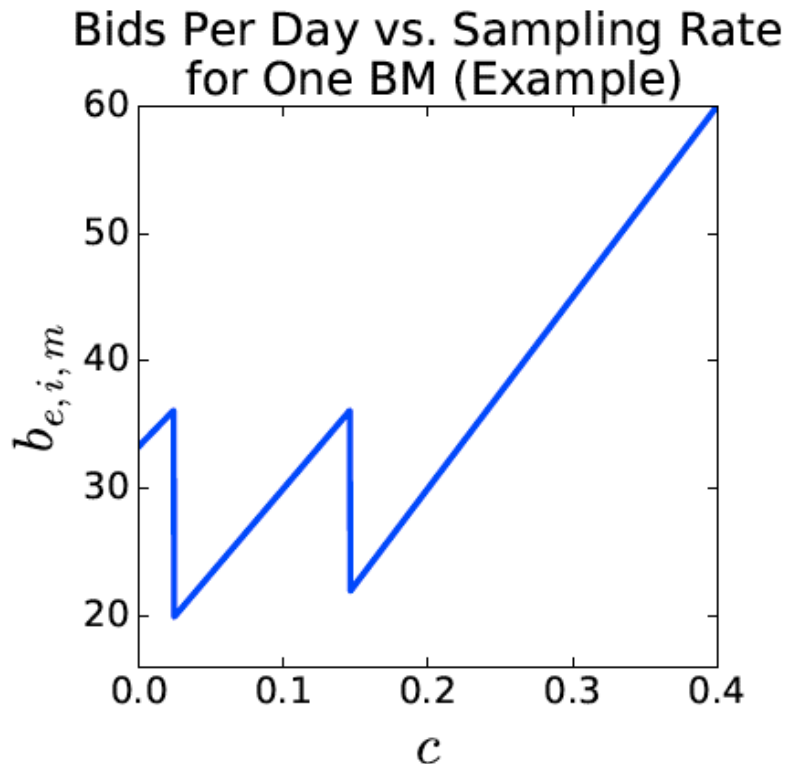


Example

Number of bids made per day vs. c for a BM with $n = 150$.

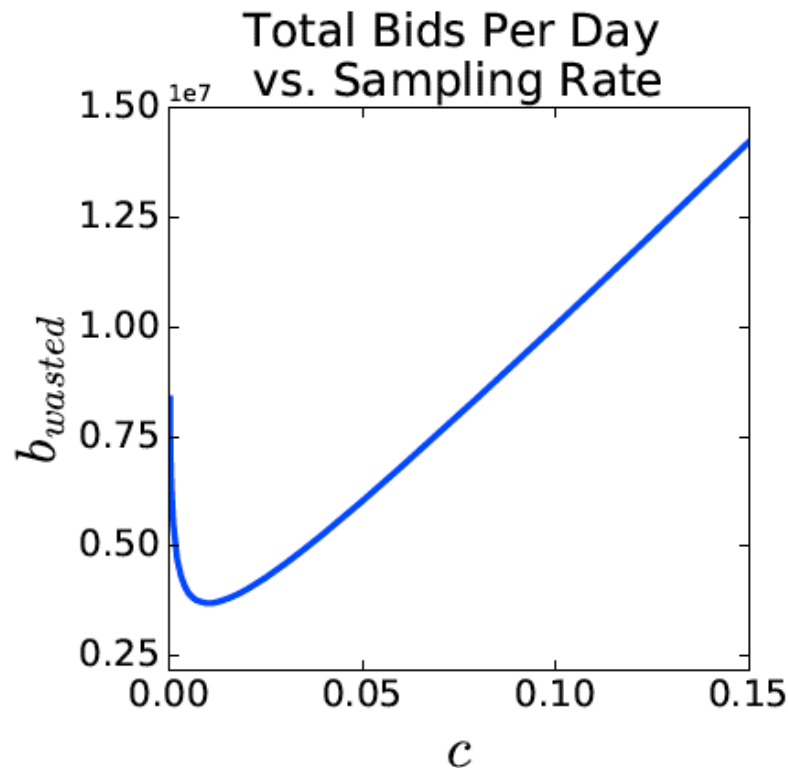
The BM combination transitions from (right to left) staying in the BM list every day to leaving the BM list for 1 out of every $T+1$ days to leaving the BM list for 2 out of every $T+1$ days.

The optimal c for this BM combination is 2%.



Optimal c for All BMs In Our System

- Optimal c for each BM depends on its n . The single optimal c for all BMs depends on the distribution of n .
- To compute the optimal c for all BMs:
 1. Identify all BMs from data over a long period of time.
 2. Minimize the sum of number of bids made to all combinations.
- The system optimal c turns out to be 1% .





Thank You

