

MEOW: A Space-Efficient Nonparametric Bid Shading Algorithm

Wei Zhang
Tsinghua University
zhang-w17@mails.tsinghua.edu.cn

Brendan Kitts
Verizon Media
brendan.kitts@verizonmedia.com

YanJun Han
Stanford University
yanjun@stanford.edu

Zhengyuan Zhou
Stanford University
zyzhou@stanford.edu

Tingyu Mao
Verizon Media
tingyu.mao@verizonmedia.com

Hao He
Verizon Media
hao.he@verizonmedia.com

Shengjun Pan
Verizon Media
alanpan@verizonmedia.com

Aaron Flores
Verizon Media
aaron.flores@verizonmedia.com

San Gultekin
Verizon Media
sgultekin@verizonmedia.com

Tsachy Weissman
Stanford University
tsachy@stanford.edu

ABSTRACT

Bid Shading has become increasingly important in Online Advertising, with a large amount of commercial [4, 12, 13, 29] and research work [11, 20, 28] recently published. Most approaches for solving the bid shading problem involve estimating the probability of win distribution, and then maximizing surplus [28]. These generally use parametric assumptions for the distribution, and there has been some discussion as to whether Log-Normal, Gamma, Beta, or other distributions are most effective [8, 38, 41, 44]. In this paper, we show evidence that online auctions generally diverge in interesting ways from classic distributions. In particular, real auctions generally exhibit significant structure, due to the way that humans set up campaigns and inventory floor prices [16, 26]. Using these insights, we present a nonparametric method for Bid Shading which enables the exploitation of this deep structure. The algorithm has low time and space complexity, and is designed to operate within the challenging millisecond Service Level Agreements of Real-Time Bid Servers. We deploy it in one of the largest Demand Side Platforms in the United States, and show that it reliably out-performs best in class Parametric benchmarks. We conclude by suggesting some ways that the best aspects of parametric and nonparametric approaches could be combined.

CCS CONCEPTS

• **Applied computing** → **Online auctions**; • **Information systems** → **Display advertising**; • **Computing methodologies** → **Machine learning algorithms**.

KEYWORDS

online bidding, shading, auction, advertising, bid, optimization

ACM Reference Format:

Wei Zhang, Brendan Kitts, YanJun Han, Zhengyuan Zhou, Tingyu Mao, Hao He, Shengjun Pan, Aaron Flores, San Gultekin, and Tsachy Weissman. 2021. MEOW: A Space-Efficient Nonparametric Bid Shading Algorithm. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3447548.3467113>

1 INTRODUCTION

Between 2017 and 2019, the Online Advertising industry underwent a massive transformation. Prior to 2017, Display ads were sold almost exclusively on Second Price Auctions. However, by 2018, First Price Auctions had increased from 5% to 43% of all auctions [10, 22]. After Google's decision to shift in 2019, 85% of display impressions were sold via First Price.

First Price Auctions present a formidable challenge to advertisers, as they require the bidder to engage in a practice called *bid shading*. Bid shading occurs when the bidder takes the private value that they would have submitted in a second price auction, and then tries to lower their bid so that it is just above the highest competing bid - this in order to minimize their price paid whilst still winning the auction [17]. This process is beset with risk, since the other bids on the auction are unknown. If the bidder shades too little, they will overpay; If they shade too much, they will lose and gain no value. Identifying the optimum bid shade, therefore, requires the bidder to *predict* competing bidder prices, yet without being able to see bidder prices. This is an enormously difficult data mining prediction problem.

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 '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467113>

With its sudden financial impact in the Online Advertising industry, Bid Shading has become a major area of new research and commercial activity. Several companies announced new bid shading services to help advertisers effectively bid, including Google [12, 13], AppNexus [4] and Rubicon [29]. Researchers and industry practitioners have also published details of new bid shading algorithms [11, 20, 28].

Most approaches for solving the problem involve predicting the probability of winning at different bid prices, the surplus given this probability of win, and then returning the shaded bid price with the maximum surplus [28]. This has generally been accomplished with parametric assumptions on the shape of the landscape.

We show new evidence in this paper that online auctions diverge significantly from well-known distributions, and exhibit usable structure due to the way that humans set up campaigns and inventory floor prices [16, 26]. Parametric distributions often fail to capture this structure.

The current paper presents a new nonparametric method for Bid Shading which enables the exploitation of this deep structure. The paper is organized as follows: Section 1 introduces the Bid Shading Problem and describes prior work including several parametric approaches that have been recently published. Section 2 introduces real auction data, and shows that it exhibits significant "spike" structure that not unlike the Left Digit Anchoring Effect observed in Psychology [16]. Section 4 introduces the MEOW algorithm, and proves its time and space complexity. Section 6 shows Offline and Live Bid Server experiments using the new algorithm.

1.1 The Bid Shading Problem

The Bid Shading Problem aims to find how much the advertiser should discount their private valuation when bidding in a First Price Auction. Let V_t represent how much the advertiser expects to capture from an impression at time t . Assuming that the valuation V_t and bid b_t are both expressed in dollars, the advertiser's financial gain, or *surplus* over a horizon T , is equal to:

$$S_T \stackrel{\text{def}}{=} \sum_{t=1}^T (V_t - b_t) \mathbf{I}(b_t > m_t), \quad (1)$$

where b_t is the shaded bid, m_t is the minimum bid price to win, and $\mathbf{I}(b_t > m_t) = 1$ if the impression is won, and 0 otherwise. The task is to find a set of shaded bids $\{b_t\}_{t=1}^T$ that maximize the surplus S_T .

1.2 Related Work

Bid shading is a common tactic in repeated First Price Auctions, and is expected by auction theory. [45] found robust evidence of shading in Austrian livestock auctions, [7] reported shading in a Texas cattle market, and [17] found the practice in auctions for US Treasury notes. Bid Shading shares characteristics with the Seller's (Reserve) Price Optimization Problem [6, 21, 23, 24, 30, 31], although the Bid Shading Problem is a buyer side problem, and the buyer has access to distinct forms of feedback. maximize profit. However buyers are trying to *minimize* their offer and receive reward in proportion to the lowness of their price, where-as sellers are trying to *maximize* their reserve price. A variety of approaches have been proposed to solve the Bid Shading problem:

Winning price predictors There have been some published research on the problem of winning price prediction on auctions. [39, 40] both develop methods for this purpose. Whilst these methods are useful, Bid shading involves predicting and maximizing expected surplus, however, which involves another unknown and optimization step.

Point Estimators use a machine learning algorithm to predict the exact optimal shading factor, by estimating the ratio of the minimum bid to win over the advertiser's private value. For instance, [11] used a Factorization Machine to predict a shading factor by training against known cases of optimum shading factor (0.1).

Unfortunately, this technique is only feasible on Seller auctions in which the Seller provides the exact winning price back to the Buyer (an "Open Bid Auction"). The optimal price can then be used as a training signal for the Buyer. However, most Online Advertising Auction Sellers (including Index Exchange, Pubmatic and others) do not disclose this information, instead just providing whether the bid was accepted or not ("Sealed Bid Auctions"). The "optimum" bid is therefore unknown.

Distribution estimators improve upon these earlier approaches, by training on the 0-1 win/loss signal, and predicting the probability of win across all possible bid prices; effectively creating a probability distribution of winning prices. Once the cdf is predicted, it is possible to calculate the expected surplus at each bid price, and the optimum bid price can be identified.

The recent WinRate model from [28] takes this approach, estimating the win probability for each bid price using a 0-1 logistic distribution, and then maximizing the resulting surplus function. Because of the known parametric form, the authors are able to bound the optimum and use a guaranteed $O(\log(K))$ bisection search to find the surplus maximum - the logarithmic time search being highly desirable for Bidding Servers which need to minimize computational operations.

Other authors have taken a similar approach to distribution estimation. [44] extend the Winrate idea to support a range of parametric distribution, and use a Deep Neural Network to estimate the distribution's shape parameters. Their implementation was shown to work for Gamma, Gaussian and Log-Normal distributions. They used a distribution-agnostic, Golden Section Search, to find the surplus maximum bid price.

Although these methods have proven effective, we show in this paper that the actual auction distributions are generally quite divergent from the parametric assumptions, and that more surplus can be captured by modeling the auction data more closely.

2 AUCTION LANDSCAPES

In auction literature and past work, Auction State bids are often assumed to be normally distributed [38, 41]; although others have noted that their auction data was fit well by Log-Normal [8]. We tested a variety of distributions on our auction data. None of the distributions fit well enough to be significant under a test for fit, but similar to [41] we found that Log-Normal has the lowest error and highest linear correlation to the actual data, out of Normal, Log-Normal and Gamma distributions tested. Log-Normal is seen in other auctions such as contract bidding [32].

In addition to the overall shape, Auction State also has some unusual characteristics. One that stands out is that there are some prices where there are large spikes in impressions. For example, Figure 1 is typical.

Some previous authors address the spikes by smoothing over them. [42] fit smoothing splines to the bid-volume cumulative distribution. [43] use a Mixture of Gaussians. If the spikes are noise, then avoiding transient spikes should result in better out-of-sample prediction. But what if the spikes are not actually noise?

Spikes appear in other domains: Marathon running times are approximately Log-normal. However, finishing times also have spikes on the left-side of hours and half hours. This is due to a Psychological effect where humans favor finishing before whole hours and half hours [2, 36].

Spikes also occur in pricing. A histogram of retail prices usually shows spikes at price ending in 9 and 5. This effect is known as the “Left Digit Anchoring Effect”, a Psychological phenomenon where consumers seem to ignore the least significant right-hand-side digits when doing value comparisons [16]. A range of theories have been offered to explain the practice, including cognitive workload from rounding up, precision being taken as an indicator of truthfulness, and others. These unusual price-points are a robust part of retail price optimization data.

In the field of online advertising, website owners set floor prices for inventory, which, in turn, impact the auction prices we observe. We encounter something unusual: The floors primarily use round numbered prices, including 5, 9.50, 10, 15, 20, 25, 30, 35, 40, 45. This can be seen clearly in Figure 2. It looks like Supply managers are susceptible to an “End Digit Effect” also!

Table 1 shows a comparison of End Digit Effects in other domains. It therefore matters whether a price of \$10 or \$9.99 is submitted to the auction – there really are “cliffs” in terms of impressions at different round price thresholds.

In order to capture these human-engineered artifacts, we need more freedom to model the data. Jacob Wolfowitz introduced the term *nonparametric*, in 1942, as a way of describing methods that did not rely on data belonging to any particular parametric family of probability distributions [37]. Estimating the auction surplus cdf by discretizing and estimating regions separately [1, 25, 34] might offer a way of modeling the unusual structure in auctions.

Some authors have tried similar approaches in the past, usually for trying to identify the hidden private valuations of bidders [14], [33], [3]. Sellers could use this information to optimize their reserve prices. However, in our case, the bidder’s private value is known, and the unknown is the clearing bid prices. The current work is the first known application of bid shading using nonparametric estimates of the auction landscape.

3 PROBLEMS WITH NONPARAMETRIC ALGORITHMS

Nonparametric methods present a range of challenges which need to be resolved before they can be used. One challenge is high storage cost. A parametric model of an auction surplus cdf will only require $O(p)$ parameters, where p is the number of features being used. A discrete approximation of the same cdf, will require $O(p \times K \times M)$ where K is the number of possible bids and M is the number of

Pennies	Online Auction Bids	Internet prices [2]	Retail prices [3]
00..09	27%	28%	8%
10..19	9%	3%	0%
20..29	8%	4%	0%
30..39	8%	4%	1%
40..49	7%	4%	0%
50..59	18%	4%	29%
60..69	6%	3%	0%
70..79	6%	4%	0%
80..89	6%	6%	1%
90..99	6%	41%	61%

Table 1: End Digit distributions. “Online Auction Bids” are from observed Verizon SSP ad prices




Figure 1: Online auctions exhibit unusual distributions of bid prices. In the above example, spikes occur at certain bid prices. This results in Normal, Log-Normal and Gamma distributions providing poor fits to this data. Furthermore attempts to smooth the spikes actually decreases predictive power. The spikes appear to be real phenomena due to human pricing effects.

private valuations. If K and M are in units of CPM with the smallest unit of bid a penny, and span all 2 place numeric values greater than 0; for bids between \$0 and \$10, and valuations between \$0 and \$100; that means $1,000 \times 10,000 = 10$ million bins.

A second challenge is generalization. Figure 3 shows performance of a fixed-width nonparametric algorithm, versus the average historical winning price. The nonparametric algorithm, in this example, works best when there are more than 80,000 auction observations. However, below this threshold, the nonparametric approach actually *performs worse* than the simple strategy of predicting the mean. The reason for this loss of performance, is because the nonparametric algorithm’s binning is too fine-grained, resulting in sparse data which doesn’t carry statistical significance. Thus, the nonparametric algorithm needs the ability to dynamically adjust its bin sizes, so that it maintains usable resolution. The MEOW algorithm, described below, uses dynamic binning.




Figure 2: Marginal distribution (bottom) versus cumulative distribution (top) spikes in price can be clearly seen at 5, 9.5, 10, 15, 20, 25, 30, 35, 40, 45. These price spikes seem to be related to the Psychology of price setting on fixed price contracts.




Figure 3: Fixed width nonparametric predictors can perform worse than the mean if their resolution is too high for the data.

4 MEOW ALGORITHM

To mitigate the issues of storage cost and generalization, we introduce a nonparametric algorithm called the Multi-resolutional Exponential Weighting (ME(O)W) algorithm. This algorithm inherits the nice theoretical properties of the Exponential Weighting based algorithm [15] through the lens of online learning, and uses a

valuation bid	v_1	v_2	$\cdots$	v_j	$\cdots$	v_M
b_1				$S_{1,j}$		
b_2				$S_{2,j}$		
$\vdots$				$\vdots$		
b_i				$S_{i,j}$		
$\vdots$				$\vdots$		
b_K				$S_{K,j}$		

Table 2: A reward table used in the SEW algorithm.

dynamic and data-driven binning to significantly reduce the memory requirement and adapt to different natures of data.

4.1 Algorithm Overview

The MEOW algorithm is motivated by the general idea of exponential weighting in nonparametric bid shading, where both the private values and bidding prices are quantized into discrete levels, and we maintain a table of historic rewards with each entry corresponding to a given pair of private value and bidding price. At each time, the private value is computed, and the bidder's bid is determined by running an exponential weighting algorithm on the rewards of all candidate prices given this private value. An example of the reward table is illustrated in Table 2, where $S_{i,j}$ denotes the cumulative historic surplus of bidding b_i under the private valuation v_j . Under the private value v_j , the exponential weighting algorithm selects a random bid b_i with probability

$$p_i = \frac{\exp(\eta S_{i,j})}{\sum_{k=1}^K \exp(\eta S_{k,j})},$$

where $\eta > 0$ is a properly chosen learning rate.

However, maintaining such a static table is typically very memory-consuming, leaving lots of bins seldom visited, and a large portion of candidate prices probably too bad for the bidder to bid. Also, the non-data-driven nature of the table leads to a poor generalization performance. The MEOW algorithm improves over the static table by choosing its rows and columns in a dynamic and data-driven way, and specifically greatly reduces the quantization levels for both the private value (horizontally) and bidding prices (vertically).

Horizontal: private value bins. The high-level idea of horizontal binning is to adapt the bin design to the real data distribution, where each bin has comparatively similar amounts of data. Specifically, if some bin of private values consists of too much data, we further split it into smaller bins to reduce the quantization error. On the other hand, if some bin has too little data, we merge it into another bin so that there is enough data in this bin for learning. In the MEOW algorithm, we first fix a static binning, and then perform the splitting and merging operations of bins based on incoming data. To reduce the computational cost, in the algorithm these steps are only performed every T_1 rounds of auctions, where $T_1 > 0$ is a hyperparameter which is moderately large (e.g. $T_1 = 1,000$).

Vertical: bidding price levels. The redundancy in the price levels comes from the fact that, the optimal bidding price given a private

value in a small bin also lies in a small range. Therefore, we could roughly estimate the optimal price (possibly with a low precision) and then keep only a few candidate prices around it. Specifically, for each bin of private values, we use the historic data to compute the empirically optimal bidding price p^* in this bin, and the set of candidate bidding prices is chosen to be a suitable quantization of $[p^* - \Delta, p^* + \Delta]$ for some small $\Delta > 0$. The final quantization level could be as small as $5 \sim 20$, which greatly reduces the storage cost. In the MEOW algorithm, the process of updating candidate bidding prices is implemented every T_2 rounds, where $T_2 > 0$ is a relatively long time (e.g. 1 day).

Discount factor. The final MEOW algorithm also involves a discount factor $\sigma \in (0, 1)$ for two purposes. First, practical data are typically non-stationary over time, and gradually forgetting old data enables a better adaptation to the new data. Second, as the amount of data increases, without the discount factor it is possible to have infinite bins of private values, which increases both the computational and the storage cost. In contrast, with a discount factor, the number of bins is always bounded from above (cf. Theorem 4.1). In the MEOW algorithm, we will apply this discount factor to both the data counts of each bin, and also the cumulative reward of each candidate bidding price.

4.2 Algorithm Details

The complete description of the MEOW algorithm is summarized in Algorithm 1, which also takes Algorithm 2 as a subroutine which updates the candidate bidding price every T_2 time steps. Specifically, the MEOW algorithm maintains an array of private value bins, where each bin is a data structure consisting of the following variables:

- $[v_{\text{low}}, v_{\text{high}}]$: range of the private value in the bin;
- count: cumulative (discounted) amount of past data falling into this bin;
- price $[K]$: an array of K candidate bidding prices under this bin, sorted in an increasing order;
- history $[K]$: an array of K cumulative (discounted) historic rewards associated with the above K bidding prices.

where $K > 0$ is a fixed parameter in the algorithm and denotes the number of vertical quantization levels.

In the initialization of the algorithm, we uniformly partition the interval $[0, V]$ into M_0 bins, where $V > 0$ is an upper bound for most private values (e.g. the 1% quantile). For each bin, we initialize K price levels to be a uniform quantization of $[0, P]$, where $P > 0$ is the maximum bidding price. All the counts and the reward histories are initialized to be zero.

Next we describe the dynamic updates of the private values and candidate prices, respectively. For the private value bins, if some private value above V occurs (which is unlikely), we create a new bin for this value. After every T_1 time steps, we check the size of each bin: if the bin count is larger than a threshold N_1 , we split it evenly into two bins, with both the count and history halved; if the bin count is smaller than another threshold N_2 , we merge it with one of its neighboring bin which a smaller size, combine their counts, and inherit the price levels and history from the larger bin. We repeat this process until the count of each bin is between

Algorithm 1: Multi-resolutional Exponential Weighting (MEOW)

Inputs: Initial number of bins M_0 ; Initial ranges V, P ; Number of prices K ; Discount factor $\sigma \in (0, 1)$; Learning rate $\eta > 0$; Update periods T_1, T_2 ; Thresholds N_1, N_2 .
Initialization: Build M_0 bins equally for $v_t \in [0, V]$, and set $\text{bin.price}[j] = jP/K$ for each $j = 1, \dots, K$.
for $t = 1, 2, \dots$ **do**
 % Search for current bin
 Observe private value v_t ;
 if $v_t > V$ **then**
 Create a new bin $[\text{floor}(v_t), \text{floor}(v_t) + 1)$;
 end
 Search for the bin* s.t. $v_t \in [\text{bin}^*.v_{\text{low}}, \text{bin}^*.v_{\text{high}}]$;
 % Exponential weighting
 for $j = 1, 2, \dots, K$ **do**
 $\text{prob}[j] \leftarrow \exp(\eta \cdot \text{bin}^*.\text{history}[j])$
 end
 Sample $b_t \sim \text{prob} / \sum_{j=1}^K \text{prob}[j]$;
 % Bin update
 Observe the minimum-bid-to-win m_t ;
 for $j = 1, 2, \dots, K$ **do**
 $\text{bin}^*.\text{history}[j] \leftarrow \text{bin}^*.\text{history}[j] + \text{instantreward}(\text{bin}^*.\text{price}[j]; v_t, m_t)$;
 end
 $\text{bin}^*.\text{count} \leftarrow \text{bin}^*.\text{count} + 1$;
 % Split or merge bins after every T_1 steps
 if $t \% T_1 == 0$ **then**
 for all bins do
 % Split a large bin into two smaller bins
 if $\text{bin}.\text{count} \geq N_1$ **then**
 $\text{bin}.\text{history} \leftarrow \text{bin}.\text{history} / 2$;
 $\text{bin}.\text{count} \leftarrow \text{bin}.\text{count} / 2$;
 Create new bins $\text{bin}_l, \text{bin}_r \leftarrow \text{bin}$;
 $\text{bin}_l.v_{\text{high}} \leftarrow (\text{bin}.v_{\text{low}} + \text{bin}.v_{\text{high}}) / 2$;
 $\text{bin}_r.v_{\text{low}} \leftarrow \text{bin}_l.v_{\text{high}}$;
 Replace bin by bin_l and bin_r ;
 end
 % Merge two smaller bins into a large bin
 if $\text{bin}.\text{count} \leq N_2$ **then**
 Find the neighbor bin' with a smaller count;
 Create a new bin* with private value range
 $[\text{bin}.v_{\text{low}}, \text{bin}.v_{\text{high}}] \cup [\text{bin}'.v_{\text{low}}, \text{bin}'.v_{\text{high}}]$;
 $\text{bin}^*.\text{count} \leftarrow \text{bin}.\text{count} + \text{bin}'.\text{count}$;
 $\text{bin}^*.\text{price}$ and $\text{bin}^*.\text{history}$ inherit from the
 bin with a larger count;
 Remove bin and bin', and add bin*;
 end
 % Discount
 Multiply all counts and histories by the factor σ ;
 end
 % Update price levels after every T_2 steps
 if $t \% T_2 == 0$ **then**
 Requantization();
 end
 end

Algorithm 2: Requantization

Global inputs: private value bins, number of price levels K

for all possible bin do

$j^* \leftarrow \text{argmax}(\text{bin.history});$

for $j = 1, 2, \dots, K$ **do**

$\text{bin.price}[j] \leftarrow \text{bin.price}[j^* - 7] + j \cdot$

$(\text{bin.price}[j^* + 7] - \text{bin.price}[j^* - 7])/K;$

end

$\text{bin.history} \leftarrow 0;$

end

$[N_2, N_1]$, and then apply the discount factor $\sigma \in (0, 1)$ to both the counts and the historic rewards.

As for the updates of candidate prices, we call the subroutine in Algorithm 2 every T_2 time steps. In Algorithm 2, for each bin we pick the best price p^* giving the largest historic reward, and update the new prices to a uniform K -level quantization of $[p^* - \Delta, p^* + \Delta]$ for some Δ . The specific choice of the interval is based on the past price level: the algorithm finds the best bidding price $\text{bid.price}[j^*]$, and chooses the interval to be $[\text{bid.price}[j^* - 7], \text{bid.price}[j^* + 7]]$. Finally, since the price levels have changed, we also reset all historic rewards to zero.

Finally we provide an example choice of the hyperparameters used in our experiments: $M_0 = 40, V = 100, P = 10, K = 20, \sigma = 0.99, \eta = 1, (T_1, T_2) = (1000, 1 \text{ day}), (N_1, N_2) = (2500, 10000)$. Note that we will ignore the first few T_1 steps to ensure enough data for splitting/merging when we restart the algorithm.

4.3 Time and Space Complexity

In this section we provide the analysis on the space and time complexities of the MEOW algorithm, and show that it could indeed be efficiently implemented in practice. We start by showing that thanks to the discount factor, the number of bins is always finite.

THEOREM 4.1. *Even for an infinite amount of data, the total number of bins is upper bounded by a constant number*

$$M := \max \left\{ \frac{T_1}{N_2 \cdot (1 - \sigma)}, M_0 \right\}.$$

PROOF. First we show that due to the discount, the total count is bounded by a constant value:

$$\sum_{\text{all possible bins}} (\text{bin.count}) \leq T_1 \cdot (1 + \sigma + \sigma^2 + \dots) < \frac{T_1}{1 - \sigma}.$$

Since after each horizontal bin update, the count of each bin is at least N_2 . In view of the above inequality, the number of bins after update is at most $T_1/(N_2(1 - \sigma))$. Moreover, before all bin updates the number of bins is initialized to be M_0 , and the result follows. $\square$

Under the choice of parameters $T_1 = 1000, N_2 = 2000, \sigma = 0.99$, and $M_0 = 40$, we compute that $M = 40$ in Theorem 4.1. Consequently, our storage cost is at most $O(MK)$, corresponding to the storage of the matrix consisting of all historic rewards.

As for the computational complexity, note that whenever there is no horizontal or vertical update, the running time of the bin search and the exponential weighted prediction in Algorithm 1 is at

most $O(\log M + K)$. When there is either a horizontal or a vertical update, we may need to change the history table for all bins, which takes $O(MK)$ time. Therefore, the overall time complexity during T rounds of auctions is

$$O \left(T \cdot \left(\log M + K + \frac{MK}{T_1} + \frac{MK}{T_2} \right) \right),$$

which is linear in T with the coefficient smaller than 30 under our parameter configuration.

5 IMPLEMENTATION

The bid shading system was deployed on Verizon Demand Side Platform (VZDSP) [27], one of the largest Real-Time Bidding systems in the United States after Google, Amazon, and the Trade Desk [9]. The performance requirements for VZDSP are extreme. At run-time, the Bid Shader needs to respond to 5.5 million requests per second peak load. For each of these bid requests, a bid needs to be calculated within 100 milliseconds. Approximately 1000 bid servers are used to serve ads, which means that each server has to handle 5,000 requests per second. Overall, less than 10 milliseconds are budgeted for bid calculations. The time complexity of Section 4.3 shows that the algorithm only adds about 30 additional operations per request.

Space requirements are also highly restrictive. Bid servers carry about 28 Gigabytes of RAM. There are over 200,000 sub-domains and mobile applications requesting bids. Therefore assuming $M=40$ and $K=20$, there are $200,000 \times M \times K = 160,000,000 = 160$ million double types are needed, which equals about 1.28 Gigabytes RAM per bid-server. The analysis to follow shows results for the most frequent 100 domains, the memory consumption for which was negligible at 0.64 Megabytes.

6 EXPERIMENTS

In order to measure the performance of the bid shading system, two forms of testing were performed: (i) previously saved auctions were replayed using the MEOW algorithm to calculate new bid prices (Offline test). (ii) the MEOW algorithm was also used in a live bid server in an A/B test versus the production algorithm (Online test). The production algorithm benchmark in both cases was an implementation of log-normal distribution-based shading [28].

6.1 Offline Auction Replay

The nonparametric algorithm was first tested on saved auction data captured from the Verizon Demand Side Platform where for each auction/bid request, private valuations and highest competing bids were known. Bid requests from the top 100 top-level domains were used (cnn.com, espn.com, buzzfeed.com and other sites), and all auctions from December 22 to January 12 2021 in which the Production algorithm responded with a bid, were used. This comprised approximately 6.2 billion requests.

The data spanned an interesting period of time, since it ranged from the 2020 Christmas shopping season with high advertising prices around \$0.97 CPM, to January 2021 in which advertising prices dropped to just \$0.77 CPM.

Figure 4 shows the behavior across this period; after good performance from December 22 to 29, there's a big drop centered on January 1, 2021. Surplus yield worsens by 15% due to the new

Metric	Distr.	MEOW
surplus % of opt	47.91%	53.42%
imps % of opt	48.01%	61.21%
spend % of opt	43.03%	62.28%

Table 3: Optimality




Figure 4: Surplus captured by nonparametric algorithm, as a percent of optimal surplus from December 22 to January 12, 2021.

auctions no longer matching historical data. On January 2nd, the surplus recovers. Quantitatively, MEOW appears to have responded better to the change in distribution over the dataset. The R^2 between algorithm bid and optimal bid for Distribution was 0.903 where-as for MEOW it was 0.953. The mean absolute difference in CPM was \$0.88 and \$0.64; and surplus as a percent of optimal surplus was 47.9% and 53.4% for Distribution vs MEOW (Table 3). Thus MEOW submitted bids that were closer to optimum and had a better correlation in matching the in-time changes to the optimum bid distribution.

Table 4 summarizes the performance on replay data. The increase in surplus ranged between 5.7% (100th percentile) to 10.1% (90th percentile), and all increases were statistically significant ($p < 0.01$; paired t-test; MEOW vs Distribution surplus scores compared daily). The table reports on several surplus percentiles because we have found it to be common for a tiny percentage of advertiser ads to have unrealistic goals and be "chronically wound up" by the control system, resulting in spuriously high surpluses. Therefore we presented a range of percentiles from 90% to 100% to help verify that the results were robust.

6.2 Online Production Performance

The algorithm was also deployed in the Verizon DSP bid server [27], and was set to run on a randomly selected 1 percent of traffic, and the top 100 domains. The period of data analyzed spans from January 21 to January 29 2021.

The results are shown in Table 5. The surplus increased between 3.3% to 6% (90th..100th percentile; all except the 100th percentile were significant at $p < 0.01$; paired t-test).




Figure 5: Predicted Optimal Shading Factor by Private Value average over all domains.






Figure 6: Predicted Surplus given Private Value and Shading Factor for one domain.

Figure 7: Comparison of probability of win for (A) Point Estimator [11], (B) Distribution Estimator [28], (C) MEOW with $K=30$ bins, and (D) MEOW with $K=7$ bins.

percentile	90	95	98	99	100
mean	10.1%	11.5%	10.4%	8.7%	5.7%
stdev	4.4%	4.1%	3.6%	3.3%	1.8%
stderr	1.0%	0.9%	0.8%	0.7%	0.4%
ttest	<0.001	<0.001	<0.001	<0.001	<0.001

Table 4: Surplus Offline Results

percentile	90	95	98	99	100
mean	3.3%	6.1%	4.5%	3.5%	6.0%
stdev	2.7%	3.3%	4.8%	3.9%	22.0%
stderr	0.72%	0.89%	1.28%	1.03%	5.87%
ttest	<0.001	<0.001	0.004	0.003	0.22

Table 5: Surplus Online Results**Figure 8: Offline MEOW versus Distribution Bid Shading algorithm Dec 22 - Jan 12 2021. Bid prices decrease significantly starting January 1**

6.3 Observations

Some examples of MEOW behavior on real Online Advertising auction data (Offline and Online experiments) are shown in Figures 5, 6, 7.

Figure 7 shows the performance MEOW approximating the actual win cumulative probability distribution for ebay.com in its offline experiment. For each bid price, the MEOW algorithm submits its predicted optimal bid. We have aggregated those bid submissions into a cumulative probability distribution. We then compared this against the actual win probability distribution in the data.

The Distribution based approach [28] does a good job of approximating the optimal bid price distribution. However there are clear imperfections in high bids; indeed we sometimes find that the Distribution estimates are biased systematically due to the shape that they are required to fit. We also observe that the Distribution approach has difficulty setting the win probability to zero for the "floor" price of the auction - instead of an immediate drop to zero, it is a gentle slope.

For illustration purposes, we show a "low resolution" MEOW algorithm that only has $K=7$ bins; the algorithm approximation is relatively poor. In contrast, the MEOW algorithm with $K=30$ bins approximates the distribution extremely well - and better than the

Distribution approach. In particular, the nonparametric approach approximates the floor (bid prices below \$2.00) and the ceiling.

Figure 5 shows the importance of private value quantization. This shows the average shading factor (bid price submitted as a percent of private value) which maximizes surplus, at each quantized private value. As private value increases, the algorithm finds that a deeper shading factor is optimal, a result also observed in the auction literature [5].

Figure 6 shows the inferred relationship between private value, shading factor, and surplus, for one domain (spotify.com). The surplus maximum "ridge line" is visible. It can also be seen that there is a bid region where the system has zero probability of winning; this is likely the inferred auction floor.

7 DISCUSSION

The higher yield from nonparametric algorithms isn't free. Whereas the Distribution algorithm might typically have two numeric parameters for its distribution shape (variance and mean for example), MEOW has $M \times K$ parameters; which for the default parameters of $M=40$ and $K=20$ results in 800 doubles. Thus, the algorithm is about 400 times more expensive in space. We've argued, and the experimental results also support, the argument that this additional space is needed to capture the various spike patterns. However, the same level of resourcing might not be necessary for every auction.

We believe it might be possible to combine both parametric and nonparametric approaches, and use the higher precision of nonparametric where needed, and preserve storage when parametric approximates well enough. One approach that seems promising is the *online learning with hints* framework from [35], where the parametric fit forms a hint which is used or discarded based on performance.

8 CONCLUSION

The shift to First Price has been traumatic for the advertising industry. Several researchers reported that traffic prices for First Price Auctions increased between 5% and 50% higher compared to Second Price Auctions [4, 18, 22, 29], meaning significantly lower advertiser profitability for the same impressions. [22] reported that after their SSP switched to First Price, 10% of advertisers actually discontinued bidding.

As a result of these problems, there has been an explosion of research and commercial implementations in Machine Learning for Bid Shading. It seems certain that the financial imperative to shade better than competing Demand Side companies, will lead bidders to begin to exploit the deep pricing structure in online auctions. There appears to be plenty of performance available, for researchers and companies who are willing to "listen to what their data is telling them".

As Thomas Huxley, the great biologist and supporter of Charles Darwin suggested, we should endeavor to *"sit down before fact as a little child, be prepared to give up every preconceived notion, follow humbly wherever and to whatever abysses nature leads, or you shall learn nothing."* - Thomas Huxley (1860), [19]

REFERENCES

- [1] 2021. Nonparametric statistics. *Wikipedia website* (2021). https://en.wikipedia.org/wiki/Nonparametric_statistics

- [2] Eric J. Allen, Patricia M. Dechow, Devin G. Pope, and George Wu. 2014. Reference-Dependent Preferences: Evidence from Marathon Runner, Management Science. *Institute for Operations Research and the Management Sciences (INFORMS)* (2014). http://faculty.chicagobooth.edu/devin.pope/research/pdf/website_marathons.pdf
- [3] Yonghong An, Yingyao Hu, and Matthew Shum. 2010. Estimating first-price auctions with an unknown number of bidders: A misclassification approach. *Journal of Econometrics* 157, 2 (2010), 328–341.
- [4] AppNexus. 2018. Demystifying Auction Dynamics for Digital Buyers and Sellers. In *AppNexus Website*. <https://www.appnexus.com/sites/default/files/whitepapers/49344-CM-Auction-Type-Whitepaper-V9.pdf>
- [5] Pierpaolo Battigalli and Marciano Siniscalchi. 2003. Rationalizable bidding in first-price auctions. *Games and Economic Behavior* 45, 1 (2003), 38–72.
- [6] A. Blum and J. D. Hartline. 2005. Near-optimal online auctions. *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms* (January 2005), 1156–1163.
- [7] John M. Crespi and Richard J. Sexton. 2005. A Multinomial logit framework to estimate bid shading in procurement auctions: Application to cattle sales in the Texas Panhandle. *Review of industrial organization* 27, 3 (2005), 253–278.
- [8] Ying Cui, Ruofei Zhang, Wei Li, and Jianchang Mao. 2011. Bid landscape forecasting in online ad exchange marketplace. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. 265–273.
- [9] eMarketer. 2020. Top 6 DSPs in the past 12 months.
- [10] Getintent. 2018. Digital Ad Impression Share Among US Supply-Side Platforms (SSPs), by Auction Type, Dec 2017 and March 2018: % of total impressions analyzed by GetIntent, April 30, 2018. <https://www.emarketer.com/content/five-charts-the-state-of-programmatic-bidding>
- [11] Djordje Gligorijevic, Tian Zhou, Bharatbhushan Shetty, Brendan Kitts, Shengjun Pan, Junwei Pan, and Aaron Flores. 2020. Bid Shading in The Brave New World of First-Price Auctions. In *The 29th ACM International Conference on Information and Knowledge Management*. CIKM'20.
- [12] Google. 2019. Real Time Bidding Protocol Protocol Buffer v.167. <https://developers.google.com/authorized-buyers/rtb/downloads/realtime-bidding-protocol>
- [13] Google. 2019. Real Time Bidding Protocol, Release Notes. <https://developers.google.com/authorized-buyers/rtb/releases#updates-2019-03-13>
- [14] Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica* 68, 3 (2000), 525–574.
- [15] Yanjun Han, Zhengyuan Zhou, Aaron Flores, Erik Ordentlich, and Tsachy Weissman. 2020. Learning to Bid Optimally and Efficiently in Adversarial First-price Auctions. *arXiv preprint arXiv:2007.04568* (2020).
- [16] J. Holdershaw, P. Gendall, and R. Garland. 1997. The Widespread Use Of Odd Pricing In The Retail Sector. *Marketing Bulletin* 8. http://marketing-bulletin.massey.ac.nz/V8/MB_V8_N1_Holdershaw.pdf
- [17] Ali Hortaçsu, Jakub Kastl, and Allen Zhang. 2018. Bid shading and bidder surplus in the us treasury auction system. *American Economic Review* 108, 1 (2018), 147–69.
- [18] B. Hovaness. 2018. Sold for more than you should have paid. In *Hearts and Science Website*. <https://www.hearts-science.com/sold-for-more-than-you-should-have-paid/>
- [19] Thomas Henry Huxley and Leonard Huxley. 1900. *The Life and Letters of Thomas Henry Huxley*. Vol. 1. Macmillan.
- [20] Niklas Karlsson and Qian Sang. 2021. Adaptive Bid Shading Optimization of First-Price Ad Inventory. In *Proceedings of the 2021 American Control Conference* (New Orleans). 4971–4978.
- [21] N. B. Keskin and A. Zeevi. 2014. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations Research* 62, 5 (2014), 1142–1167.
- [22] Brendan Kitts. 2019. Bidder Behavior after Shifting from Second to First Price Auctions in Online Advertising. http://www.appliedaisystems.com/papers/FPA_Effects33.pdf
- [23] B. Kitts and K. Hetherington-Young. 2005. Price Optimization in Grocery Stores with Cannibalistic Product Interactions. *Proceedings of the First Workshop on Data Mining Case Studies, Fifth IEEE International Conference on Data Mining (ICDM 2005)* (November 2005), 74–91. http://dataminingcasestudies.com/DMCS_WorkshopProceedings25.pdf
- [24] R. Kleinberg and T. Leighton. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. *44th Annual IEEE Symposium on Foundations of Computer Science* (2003), 594–605.
- [25] Paul Kvam and Brani Vidakovic. 2007. *Nonparametric Statistics with Applications in Science and Engineering*. John Wiley & Sons, Inc. <http://zoe.bme.gatech.edu/~bv20/isy6404/Bank/npmarginal.pdf>
- [26] Daniel Levy, Dongwon Lee, Haipeng (Allan) Chen, Robert J. Kauffman, and Mark Bergen. 2011. Price Points and Price Rigidity. *Review of Economics and Statistics* 93, Issue 4. https://www.biu.ac.il/soc/ec/d_levy/wp/pricepointsaug2008.pdf
- [27] Christina MacDonald. 2020. Verizon Media brings native marketplace, premium inventory into expanded DSP. In *Verizon Media Website*. <https://www.verizonmedia.com/press/2020/05/01/verizon-media-expanded-dsp>
- [28] Shengjun Pan, Brendan Kitts, Tian Zhou, Hao He, Bharatbhushan Shetty, Aaron Flores, Djordje Gligorijevic, Junwei Pan, Tingyu Mao, San Gultekin, and Jianlong Zhang. 2020. Bid Shading by Win-Rate Estimation and Surplus Maximization. In *The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. AdKDD'20, San Diego, California, USA.
- [29] Rubicon. 2018. Bridging the Gap to First-Price Auctions: A Buyer's Guide. In *Rubicon Website*. http://go.rubiconproject.com/rs/958-XBX-033/images/Buyers_Guide_to_First_Price_Rubicon_Project.pdf
- [30] Ilya Segal. 2003. Optimal Pricing Mechanisms with Unknown Demand. *American Economic Review* 93, 3 (June 2003), 509–529. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=297674
- [31] Hermann Simon. 1989. *Price Management*. North-Holland.
- [32] Pablo Ballesteros-Pérez & Martin Skitmore. 2017. On the distribution of bids for construction contract auctions. *Construction Management and Economics* 35, 3 (2017), 106–121.
- [33] U. Song. 2004. *Nonparametric estimation of an eBay auction model with an unknown number of bidders*. Department of Economics, University of British Columbia.
- [34] L. Wasserman. 2006. *All Things Nonparametric*. Springer Verlag.
- [35] Chen-Yu Wei, Haipeng Luo, and Alekh Agarwal. 2020. Taking a hint: How to leverage loss predictors in contextual bandits?. In *Conference on Learning Theory*. PMLR, 3583–3634.
- [36] J. Wolfers. 2014. What Good Marathons and Bad Investments Have in Common. *New York Times* (22 April 2014). <https://www.nytimes.com/2014/04/23/upshot/what-good-marathons-and-bad-investments-have-in-common.html>
- [37] J. Wolfowitz. 1942. Additive Partition Functions and a Class of Statistical Hypotheses. *Annals of Statistics* (1942), 247–279.
- [38] W.C.H. Wu, M.Y. Yeh, and M.S. Chen. 2015. Predicting winning price in real time bidding with censored data. *KDD* (2015).
- [39] Wush Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2018. Deep censored learning of the winning price in the real time bidding. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '18, London, United Kingdom, 2526–2535.
- [40] Wush Chi-Hsuan Wu, Mi-Yen Yeh, and Ming-Syan Chen. 2015. Predicting winning price in real time bidding with censored data. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '15, Sydney, 1305–1314.
- [41] Z. Xie, K. Lee, and L. Wang. 2017. Optimal Reserve Price for Online Ads Trading Based on Inventory Identification. In *Proceedings of AdKDD Workshop*.
- [42] Qixing Zheng, Niklas Karlsson, and Qian Sang. 2018. Plant Estimation and Feedback Control for Online Advertising Campaigns. In *The 12th China-Japan International Workshop on Information Technology and Control Applications (ITCA2018)*.
- [43] T. Zhou, J. Guo, and N. Karlsson. 2018. A KPIs Forecasting Scheme for Online Advertising Systems. *Technipulse* (2018).
- [44] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhushan Shetty, Brendan Kitts, Djordje Gligorijevic, Junwei Pan, San Gultekin, Tingyu Mao, Jianlong Zhang, and Aaron Flores. 2021. Efficient Deep Distribution Network for Bid Shading in First Price Auctions. *KDD '21 Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2021).
- [45] Christine Zulehner. 2009. Bidding behavior in sequential cattle auctions. *International Journal of Industrial Organization* 27, 1 (2009), 33–42.