



 Latest updates: <https://dl.acm.org/doi/10.1145/3289600.3290969>

RESEARCH-ARTICLE

## Genie: An Open Box Counterfactual Policy Estimator for Optimizing Sponsored Search Marketplace

**MURAT ALI BAYIR**, Microsoft Corporation, Redmond, WA, United States

**MINGSEN XU**, Microsoft Corporation, Redmond, WA, United States

**YAOJIA ZHU**, Microsoft Corporation, Redmond, WA, United States

**YIFAN SHI**, Microsoft Corporation, Redmond, WA, United States

**Open Access Support** provided by:

**Microsoft Corporation**



PDF Download  
3289600.3290969.pdf  
29 January 2026  
Total Citations: 7  
Total Downloads: 302

Published: 30 January 2019

[Citation in BibTeX format](#)

WSDM '19: The Twelfth ACM International Conference on Web Search and Data Mining  
February 11 - 15, 2019  
Melbourne VIC, Australia

**Conference Sponsors:**

SIGKDD

SIGMOD

SIGIR

SIGWEB

# Genie: An Open Box Counterfactual Policy Estimator for Optimizing Sponsored Search Marketplace

Murat Ali Bayir\*  
 Microsoft Bing Ads  
 Redmond, WA, USA  
 mabayir@microsoft.com

Yaojia Zhu  
 Microsoft Bing Ads  
 Redmond, WA, USA  
 yaozhu@microsoft.com

Mingsen Xu  
 Microsoft Bing Ads  
 Redmond, WA, USA  
 mingx@microsoft.com

Yifan Shi  
 Microsoft Bing Ads  
 Redmond, WA, USA  
 yifanshi@microsoft.com

## ABSTRACT

In this paper, we propose an offline counterfactual policy estimation framework called Genie to optimize Sponsored Search Marketplace. Genie employs an open box simulation engine with click calibration model to compute the KPI impact of any modification to the system. From the experimental results on Bing traffic, we showed that Genie performs better than existing observational approaches that employs randomized experiments for traffic slices that have frequent policy updates. We also show that Genie can be used to tune completely new policies efficiently without creating risky randomized experiments due to cold start problem. As time of today, Genie hosts more than 10000 optimization jobs yearly which runs more than 30 Million processing node hours of big data jobs for Bing Ads. For the last 3 years, Genie has been proven to be the one of the major platforms to optimize Bing Ads Marketplace due to its reliability under frequent policy changes and its efficiency to minimize risks in real experiments.

## CCS CONCEPTS

- Information systems → Sponsored search advertising;
- Computing methodologies → Machine learning; Modeling and simulation.

## KEYWORDS

Sponsored Search, Counterfactual Policy Estimation, Causal Inference

### ACM Reference Format:

Murat Ali Bayir, Mingsen Xu, Yaojia Zhu, and Yifan Shi. 2019. Genie: An Open Box Counterfactual Policy Estimator for Optimizing Sponsored Search Marketplace. In *The Twelfth ACM International Conference on Web Search and*

\*Corresponding author.

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.

WSDM '19, February 11–15, 2019, Melbourne, VIC, Australia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5940-5/19/02...\$15.00

<https://doi.org/10.1145/3289600.3290969>

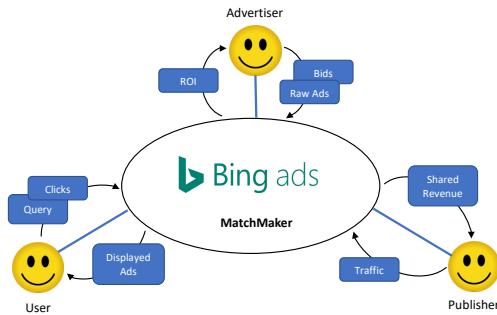


Figure 1: Sponsored Search Marketplace

*Data Mining (WSDM '19), February 11–15, 2019, Melbourne, VIC, Australia.*  
 ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3289600.3290969>

## 1 INTRODUCTION

Online advertising [10, 20] becomes very important as we spend more time on internet based applications like search engines, social network web sites and video streaming services. Sponsored search [3, 6] represents the largest share of the online advertising [20, 27] where relevant ads (advertisements) are shown to search engine users while they're looking for information. Due to high profitability and user engagement, large companies like Microsoft, Yahoo and Google make significant investments to optimize their sponsored search marketplace.

Sponsored Search Marketplace has four major actors (Figure 1) in most general form. Search engines like Bing performs match-making role between user, advertiser and publisher. Matchmaker is responsible for delivering relevant ads to the user through the publisher as well as getting bid and raw ads data from advertiser. In some cases, publisher and matchmaker can be same party like Bing. While users interact with the publisher web site, the matchmaker provides ads in back-end and shares the revenue with the publisher. When users are exposed to relevant ads related to their intention, click events lead to return of investment (ROI) for advertisers.

Optimizing Sponsored Search Marketplace is not an easy process since each actor has their own goals and these goals may conflict with each other. The goal of optimization process is to find an equilibrium point which satisfies the constraints coming from all of the parties. The most challenging part of the optimization process is the

predicting the Key Performance Indicator (KPI) impact of any new modification to the system which is also known as counterfactual policy estimation problem [1, 16].

Counterfactual policy estimation problem attracted many researcher recently within both organic and sponsored search problem context [2, 16, 22]. *A/B* testing [13, 15] is known to be most reliable solution to this problem for large scale web applications like sponsored search systems. In *A/B* testing, users are randomly divided into treatment and control groups and receive ads from two different systems. In this setup, control users receive ads from unmodified application while treatment users receive ads from the application where modification is applied. After that, KPIs from control and treatment traffic slices are compared. While providing very reliable KPI comparison, *A/B* testing is very difficult to use for many use cases due to end to end deployment requirement and traffic size limitations.

Existing observational approaches like Inverse Propensity Score Weighting [2, 8, 12, 19] are another family of counterfactual estimation methods that can be used to perform KPI prediction for policy changes. The idea is to enable subset of real traffic with randomized experiment that represents the modification to the system. Once enough data from randomized traffic is observed with the treatment modification, the offline exploration can be performed on this data to find better operating points. The offline exploration process consumes randomized input to KPI metrics mapping data without running the whole system in open box manner (system is treated as a **closed box**). While improving traffic limitations compared to *A/B* testing, existing observational approaches need end to end deployment and exploration in real traffic could be very costly depending on the size of the exploration space. In addition, observational approaches can show large regressions for the data that has bias due to frequent policy updates.

In this paper, we propose a new counterfactual estimation framework called Genie that addresses weak points of previous approaches. Genie utilizes offline simulation engine that performs replaying past event logs in an open box manner and compute KPIs via trained user model from search logs. We claim that Genie has more knowledge on inter dependencies of **system layer** that impacts the code path from the point where modification applied till to get measurement result. Therefore, replaying the historical data with proper modeling of system inputs (user and/or advertiser signals) becomes very reliable and provide many advantages compared to *A/B* Testing and existing observational approaches. Apart from KPI prediction reliability, Genie can explore much wider parameter space since it does not require real traffic with modification/exploration cost. Genie can also be leveraged to tune completely new policies where creating initial experiment is very costly due to cold start problem.

Due to its advantages, Genie has been one of the best policy estimation methods for optimizing majority of use cases in Bing Ads that includes adding new policies, large changes in existing policies and any change that is difficult/risky to validate via real traffic. As time of today, Genie hosts more than 10000 production optimization jobs yearly for Bing Ads Sponsored Search Marketplace. To the best of our knowledge, this paper is the first attempt that propose a log replay based counterfactual estimation system under proper user click behavior modeling. In particular, contributions of this paper are given below:

- We modeled Sponsored Search optimization problem as a causal inference problem to predict KPI outcome of any modification to the system. We propose a causal graph with separate input and system layers which can introduce different types of bias in the training data.
- We showed that using log replay as an open box simulator for the system layer of causal graph with the proper Machine Learning modeling of input layer yields very reliable KPI predictions.
- We also showed that Genie can easily be leveraged to tune completely new policy which is one of most risky use cases for optimizing sponsored search marketplace. We found that KPIs for the best operating points discovered by Genie is statically significant and correlated with KPIs from the real traffic when modification is applied under *A/B* testing.

This paper is organized as follows: Section 2 discusses background and motivation for this work. The next section summarizes the Related Work. Section 4 introduces the details of Genie counterfactual estimator framework. After that, Section 5 presents our experimental results. Finally, we give our conclusions in Section 6.

## 2 BACKGROUND AND MOTIVATION

In statistics, causal inference literature defines counterfactual as any change to the system the impact of which is the focus of interest [1]. In the case of sponsored search, counterfactual could be any policy, parameter or model change in the system that yields a different ad allocation presented to the end user.

Similar to the other application contexts, counterfactual estimation for the sponsored search optimization can be modeled as a causal inference problem [16]. For large systems like Bing Ads, the causal graph may contain large system layer with many active policies depending on user and advertiser signals (Figure 2).

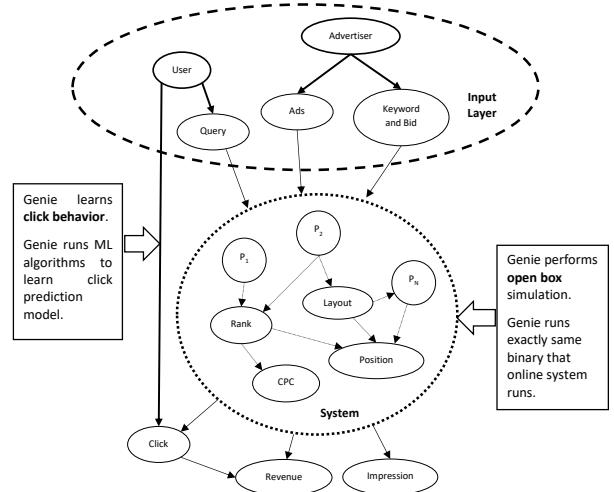


Figure 2: Causal Graph for Sponsored Search

Our main motivation for pushing simulation based counterfactual estimation system comes from the complexity of the causal graph. At a given snapshot, Sponsored Search system may have  $N$  active policies  $A = \{P_1, P_2, \dots, P_N\}$  depending on system inputs.

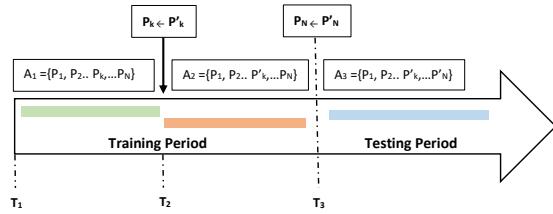


Figure 3: Example Tuning with Noise

For sophisticated systems like Bing Ads, there are frequent updates on the set of active policies (in/out) for each deployment which could potentially create a noise when historical data is used for counterfactual estimation of future KPIs. Consider the tuning process in Figure 3 that optimizes parameters of policy  $P_N$  which is the member of active policies  $A_1$  in the beginning of the training period. Let say that there is another experiment enabled in all of the traffic slices at time  $T_2$  which updates the policy  $P_k$  and changes the distribution of inputs that policy  $P_N$  consumes. In this particular case, the training time window is needed to be reset from  $T_2$  and this period should be extended beyond  $T_3$  to include enough data with the new distribution. However, this could not be practical in many cases due to the frequency of updates and timing constraints.

Unlike existing observational approaches, our simulation based counterfactual estimation can consume all of data for  $[T_1, T_3]$  time period without resetting tuning process. We propose that KPI impact of any policy can be estimated by replaying training data with the modified policy and using user click behavior model that has tolerable noise. In our policy estimation model, we keep logged advertiser data and user queries fixed while updating click signal based on new page allocation produced by running system layer with new modifications. During the tuning process, the data selection for input layer are randomized by using different seeds to prevent bias due user and advertiser signals.

For the example scenario in Figure 3, if the simulation system has an access to the logged system inputs (user and advertiser signals) and implementation of policy change (usually weeks before creating experiment) in the online binaries for  $P_k'$ , the training data can be simulated with counterfactual policy setting of  $A' = \{P_1, P_2, \dots, P_k', \dots, P'_N\}$  for the time period between  $T_1$  and  $T_3$ .

Apart from supporting tuning scenarios on noisy data given in Figure 3, our simulation platform has other advantages as follows:

- Genie could easily used for finding initial settings for randomized experiments to minimize cost of cold start problem.
- Applying modification or randomizing system inputs are not possible in certain cases (like bid of advertisers) due to legal issues or traffic capacity. Genie could easily estimate the impact of change in system inputs without real traffic.
- Randomized experiments could be very risky in certain cases for new policy changes in terms of user experience. Genie could easily be used for estimating KPI Impact of these changes before enabling in real traffic.

The reader should also be aware of that some scenarios are difficult to tune both with Genie and existing observational approaches. A typical example is where tuning setup has significant deviations

from existing policies/models in real traffic. This type of modifications yield large change in feature distributions of click prediction that Genie is using. For existing observational approaches, these type of tunings require large exploration space which can significantly regress the user experience. A hypothetical example of this category is that if traffic slice is not tuned for a long time (on the order of several months), there could be large change needed in existing operating points of many policies since marketplace data changed a lot.

### 3 RELATED WORK

Counterfactual reasoning and estimation systems [2, 16, 23] answer questions like "how would the system performance had been changed if the modification had been applied to the system during log collection?". The idea for these systems is to use online event logs to train models and predict KPIs for new modifications.

*A/B* testing [14, 15] is a standard way of evaluating the modifications to the web based systems. In *A/B* testing setup, users are randomly divided into control and treatment groups. Control group is served by existing binary without any modification whereas treatment group is served by the system where modification is applied. Once enough data is collected from *A/B* test, control and treatment metrics are compared to test the hypothesis that supports modification to the system.

Observational approaches [1] like Inverse Propensity Weight Estimation [2, 8, 12] were proposed as an alternative to the *A/B* testing for counterfactual estimation of KPI metrics. Estimation methods like Importance sampling [2] can be used in conjunction with online randomization of tunable parameters. Once the logs with randomized parameters are collected, offline exploration can be used to estimate KPI metrics for the proposed parameter settings. As mentioned in previous sections, observational approaches could not be practical in some of tuning scenarios and online explorations should be designed very carefully to minimize negative impact on user experience. [21].

Our focus in this paper is to propose efficient counterfactual estimator for practical tuning problems within the context of Sponsored Search System. In particular, we target on tuning scenarios where *A/B* testing and existing observational approaches can not work efficiently. We split the causal graph for the sponsored search problem into input and system layer and apply different methodologies to predict outcomes. While machine learning methods gives very good results for predicting user click behaviors [17, 18, 26], we used log replay based **open box simulator** to estimate the outcome of the system layer. The reader could think of applying similar approaches for policy optimization problem when training data is very noisy due to frequent system changes or when it is not practical to create online experiment with the proposed modification.

### 4 GENIE COUNTERFACTUAL ESTIMATOR

Genie Counterfactual Estimator corresponds to training job called Auction Simulation that is executed at Microsoft Cosmos Cloud [4]. Auction Simulation job takes many parameters for traffic slice filtration, click model features and user grid that contains counterfactual parameters. Depending on the size of the hyperparameter space,

Auction Simulation job can be followed by grid exploration step to improve the efficiency of solution search in hyperparameter space. In the next section, we discuss details of Auction Simulation Job. After that, Grid Exploration step is presented.

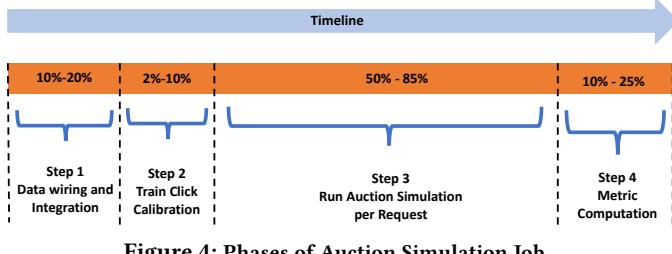


Figure 4: Phases of Auction Simulation Job

## 4.1 Auction Simulation Job

Auction Simulation Job runs Scope Script on Cosmos [4] to perform the counterfactual estimations. Figure 4 presents stages of this job including their relative minimum and maximum computation cost. The details of each phase is discussed in the next subsections.

**4.1.1 Data Wiring and Integration.** The goal of Data Wiring and Integration phase is the reconstructing online library inputs from past event logs. This phase wires logged data from multiple resources that contains ads, auction parameters and page layout data. Then, all of data sources are integrated and transformed into schemas that online library expects. The validation of logged data is performed in this step too.

**4.1.2 Train Click Prediction.** This step is responsible for training model that is used for click prediction. Click prediction is required to re-estimate the click probability of ads in the new page allocation returned from Auction Simulator when counterfactual modification is applied.

Genie click prediction model uses set of  $N$  impression feature vectors and click pairs  $D = \{(x_1, y_1), \dots, (x_k, y_k), \dots, (x_N, y_N)\}$  where  $\forall y_k \in \{-1, 1\}$  (1: click, -1: no click) as an input. The goal of the click prediction step is to find mapping from each impression vector  $x_k$  to click probability  $\hat{y}_k$  where each  $\hat{y}_k \in [0, 1]$ .

Due to high runtime performance and high prediction accuracy, majority of Genie estimation jobs are using online version of Bayesian Probit [11] for the click prediction problem.

Bayesian Probit is a generalized linear model with the following posterior distribution of  $y$  for given  $x$  and weight vector of  $w$  of linear model.

$$p(y|x, w) = \Phi\left(\frac{y \cdot w^T \cdot x}{\beta}\right) \quad (1)$$

The probit function  $\Phi(t)$  is the standardized cumulative Gaussian density function that maps the output of the linear model in  $[-\infty, \infty]$  to  $[0, 1]$  with  $\beta$  that represent steepness of probit function.

$$\Phi(x) = \int_{-\infty}^x N(x, 0, 1) dx = \int_{-\infty}^x \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx \quad (2)$$

In Bayesian Probit, each impression vector  $x_i$  contains  $L$  features. Each feature  $j$  corresponds to  $M_j$  different discrete bins and represented by a binary vector as below:

$$x_{ij} = \begin{bmatrix} x_{ij(1)} \\ \dots \\ x_{ij(M_j)} \end{bmatrix} \text{ and } \sum_{k=1}^{k=M_j} x_{ij(k)} = 1 \quad (3)$$

Since each feature  $j$  of  $x_i$  corresponds to  $M_j$  dimensional vector of bins,  $x_i$  is represented by  $\sum_{k=1}^{k=N} M_k$  dimensional sparse vector as follows:

$$x_i = \begin{bmatrix} x_{i1(1)} \\ \dots \\ x_{iL(M_L)} \end{bmatrix}, \sum_{j=1}^{j=L} \sum_{k=1}^{k=M_j} x_{ij(k)} = L \quad (4)$$

The probit model stores as a vector of gaussian distributions for each bin that represents the weights of generalized linear model as follows:

$$w = (\mu, \sigma), \mu = \begin{bmatrix} \mu_{11} \\ \dots \\ \mu_{L(M_L)} \end{bmatrix}, \sigma^2 = \begin{bmatrix} \sigma_{11}^2 \\ \dots \\ \sigma_{L(M_L)}^2 \end{bmatrix} \quad (5)$$

During the model training, each request data is converted to a set of impression and click signal  $\{-1, 1\}$  features on Cosmos. Each impression vector and click signal pair is only scanned one time in an online manner. During training data scan, the mean and variance of Gaussian distributions for matched weight bins are updated via backward message passing [11] to maximize the posterior probability in Equation 1. Once the model is trained, the inferred model is sent to downstream components of the scope script. For an evaluation of given feature vector  $x$ , the Probit model accumulates variance and mean of those matched feature bins as follow:

$$\sigma^2 = x^T \sigma^2, \mu = x^T \mu \quad (6)$$

Then, predicted click probability is calculated as  $\Phi(\mu / \sqrt{\sigma^2 + \beta^2})$  where  $\Phi$  is the standardized cumulative Gaussian density function.

Apart from online Bayesian Probit, Genie also supports Click Prediction model based on Gradient Tree Boosting [5]. Tree Boosting is a state of art method that is widely used in many real world applications including search result ranking and click through rate prediction.

Tree Boosting model implemented in Genie works with both continuous and categorical features. Therefore, it does not require costly feature engineering process to create binned feature space like in Bayesian Probit. Tree boosting model has  $L$  regression trees and each tree  $T_k$  ( $k \leq L$ ) is traversed from top to down to find a particular leaf node for a feature vector  $x$  via conditions on intermediate nodes. The weight of discovered leaf  $w_k$  is used as a weight for current tree  $T_k$ . This process is repeated for all of the trees. After that, a sigmoid function is applied to the sum of weights to create final click prediction value between  $[0, 1]$  as below:

$$y_i = \sigma\left(\sum_{k=1}^{k=N} w_k\right), \sigma(x) = \frac{1}{1 + e^{-x}} \quad (7)$$

During the training phase, each request data is converted to a set of features that is very similar to the one for Bayesian Probit. After that, the training data is scanned  $L$  times in batch manner to train tree ensemble. Inspired from Friedman et al. [7], the following log

loss objective function on prediction error is propagated in each iteration  $t$  to create new trees:

$$\text{Loss}^t = \sum_{i=1}^{i=N} \psi(y_i, \hat{y}_i^t) = \sum_{i=1}^{i=N} \left[ -y_i \log \hat{y}_i^t - (1 - y_i) \log(1 - \hat{y}_i^t) \right] \quad (8)$$

$$\hat{y}_i^t = \sigma \left( \sum_{k=1}^{k=t} w^k \right) \quad (9)$$

While constructing the single tree in iteration  $t$ , the Equation 8-9 are used to produce target value for each data sample that corresponds to the gradient of the error function given in Equation 10.

$$\bar{y}_i^t = \left[ -\frac{\partial \psi(y_i, \hat{y}_i^{t-1})}{\partial \hat{y}_i^{t-1}} \right] \quad (10)$$

After that, the split gain function is derived from least squares on target values to select feature with highest gain to create new branches in current iteration. The split process is done iteratively in greedy manner starting from feature that has highest split gain.

From the counterfactual policy estimation perspective, both methods has different trade offs. While tree boosting eliminates the process for costly feature engineering, it does not fit into traditional Map Reduce paradigm well compared to Bayesian Probit since it requires multiple disk read and write for each iteration. On the other hand, we found that Tree Boosting yields models with higher accuracy in most common tuning scenarios as mentioned in Experimental Results Section. As time of today, Genie counterfactual estimator keep both methods and user config defines the one that is going to be applied based on efficiency and quality trade offs.

**4.1.3 Auction Simulation.** This phase executes Bing Ads online libraries with the modified version of the reconstructed inputs per request. The offline replay process of the online binary with same or modified inputs is called **open box auction simulation**.

The Auction Simulaton contains a wrapper layer which is responsible for reading all of the inputs and invoking the online library. Simulation wrapper reads the counterfactual grid points, click prediction model and auction data. Then, each counterfactual grid point is converted into modifier method. The Modifier method mutates the auction data in place and returns the restorer function for auction data and object representation of the current grid point. Once the online library is called with the mutated input, the page allocation for the mutated input is generated. After that, click predict model is applied to re-evaluate click probability of ads in the page allocation. Finally, KPI metrics for current settings are computed from updated page allocation. The pseudo code for the simulation logic is given in Algorithm 1.

The key part of the Algorithm 1 is the simulation of the online library (line 6) for running an auction. There are many ways to design an online library that runs an auction [9, 24, 25] in real time. Here we only discuss the general form described in Varian et al [24] based on Generalized Second Price Auction. The subsequent auction logic described in this paper is consistent with definitions in Varian et al [24]. However, our framework can run with any type of auction mechanism as long as the interface remains same.

---

**Algorithm 1:** Simulation Algorithm

---

**Input:** Auction Data:  $A$ , Grid:  $G$ , Click Model:  $C$   
**Output:** List of (Setting, KPI) pairs as  $KPIs$

```

1  $M \leftarrow G.\text{GenerateModifiers}(A)$ 
2  $KPIs \leftarrow \{\} // \text{Initialize the output.}$ 
3 foreach  $M_i \in M$  do
4   // Modify the input and get (restorer, setting)
5    $(R_i, S_i) \leftarrow M_i.\text{Modify}(A)$ 
6    $P_i \leftarrow \text{OnlineLibrary}(A) // \text{Create } i^{th} \text{ page allocation}$ 
7    $C.\text{Predict}(P_i) // \text{Adjust click probabilities}$ 
8    $KPI_i \leftarrow \text{GetKPI}(P_i)$ 
9    $KPIs \leftarrow KPIs \cup (S_i, KPI_i)$ 
10   $R_i.\text{Restore}(A) // \text{Restore the input to original value}$ 
11 end

```

---

In general form, the auction system can get 3 sets of input data from simulation wrapper. These inputs can contain basic ad data, PClick scores and other ad specific metadata. In addition to data inputs, auction layer can also get settings for different policies and candidate page layouts to select. Following [2, 24], the rank score or the utility of  $ad_i$  can depend on bid  $b_i$  of current advertiser and ad dependent click quality score  $q_i$ . The utility  $u_i$  of  $ad_i$  can be defined as  $u_i = b_i q_i$  where  $q_i$  contains click probability  $p_i$  of current ad. Once the utility computation for each ad is done, ads can be placed on page layouts starting from most significant position to the least significant in a greedy manner. For each available position in the page layout, the winner ad and runnerup can be determined based on their rank scores. The runnerup can be used for determining pricing score for the current position which is required for calculating the cost per click (CPC) based on Generalized Second Price Auction [24]. Finally, the page allocation with the maximum utility can be selected as a best allocation.

**4.1.4 KPI Computation.** KPI computation phase computes metrics of interest for each counterfactual grid setting. Aggregated metrics are used in decision process to update the operating points of several policies in the real serving environment. In this phase, metric data for each request and grid point pair are converted into data cubes. Each data cube contains  $n$  dimensions with many datacube cells. A data cube cell contains KPI metrics that is pivoted by  $n$  dimensional key. Once the data cube for each request and grid point is created, they are aggregated recursively with sum operation over all requests to produce final metrics.

## 4.2 Grid Exploration

Some of the tuning scenarios need to explore very large hyperparameter space that contains more than 15-20 policy parameters. As a result, these tuning jobs requires significant amount of computing resources to run simulation based approach. For an efficient hyperparameter space search, the Grid Exploration step can be executed between sequential tuning jobs. In this setup, grid exploration step reads simulated data points from former job and recommend better operating points for latter jobs based on the model trained from simulated points of former job.

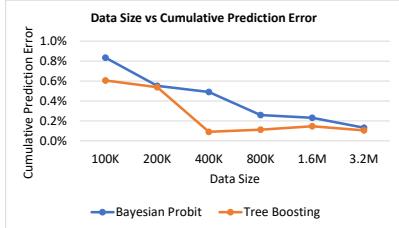


Figure 5: Cumulative Prediction Error

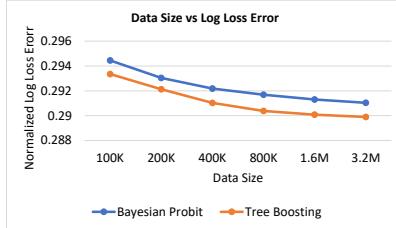


Figure 6: Log Loss Error

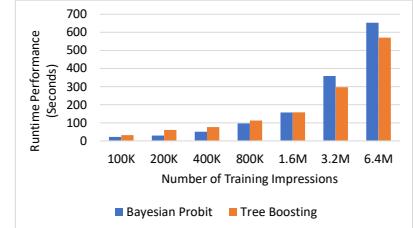


Figure 7: Runtime Comparison

In a nutshell, Grid Exploration step learns a regression model that maps set of hyperparameters  $X$  to set of real valued KPI deltas  $\Delta Y$ . KPI delta refers to normalized difference between KPI that is estimated by applying grid point (treatment) and KPI that is estimated from logs without any modification (baseline). This step learns a separate instance of regression model for each metric that tuning user is interested. For simplicity we assume that each  $\Delta y_i \in \Delta Y$  is one dimensional real value and each  $x_i \in X$  is a  $k$  dimensional vector. For a given data set  $D = \{(x_1, \Delta y_1), \dots, (x_k, \Delta y_k), \dots, (x_N, \Delta y_N)\}$  the Grid Exploration step can learn two types of regression model.

In linear regression case, the least squares approach is used to minimize the following sum which yields solution for coefficients  $\beta = (\beta_0, \beta_1, \dots, \beta_k)$ :

$$\text{Min}_{\beta} \left\{ \sum_{i=0}^N \left( \Delta y_i - \sum_{j=1}^k \beta_j x_{ij} - \beta_0 \right)^2 \right\} \quad (11)$$

For the ridge regression, the extra term is added to the error function to shrink the linear coefficients  $\beta$  with shrinkage factor  $\lambda$  to avoid potential overfitting:

$$\text{Min}_{\beta} \left\{ \sum_{i=0}^N \left( \Delta y_i - \sum_{j=1}^k \beta_j x_{ij} - \beta_0 \right)^2 + \lambda \sum_{j=0}^k \beta_j^2 \right\} \quad (12)$$

Once the model is trained, the iterative optimization algorithm similar to hill climbing is executed with the user supplied parameters. The details of the optimization process is given in Algorithm 2. In each iteration, a temporary solution set is created by exploring neighbours of the current solution set. The most common exploration logic is selecting random values from  $[min, max]$  range of random subset of each hyperparameter dimensions based on uniform distribution. Then, the best  $P$  (population size) solutions from the union of current and previous step are chosen to finalize the solution set in the current step. The iteration stops after certain number of executions and top  $k$  solutions from the solution set are selected as final result.

## 5 EXPERIMENTAL RESULTS

In this section, we provide our experimental results on Bing serving traffic. We first present experimental results on the the quality and efficiency of click prediction, auction simulation and grid exploration steps. After that, we compare the counterfactual estimation accuracy of Genie with an Importance Sampling based estimator on one of most noisy and largest traffic group in Bing. Finally, we present counterfactual estimation accuracy results of Genie on one of the most difficult use cases that requires tuning new policy.

### Algorithm 2: Optimization Algorithm

```

Input: Hyperparameters:  $X$ , KPI deltas:  $\Delta Y$ , Batches:  $B$ 
Input: Population size  $P$ , Solution size:  $k$ , Objective:  $F_{obj}$ 
Output: Best solution set:  $X'$  and  $\Delta Y'$ 
1 Model = Train( $X, Y$ )
2  $X' \leftarrow X, \Delta Y' \leftarrow \Delta Y$ 
3 for  $i \leftarrow 1$  to  $B$  do
4    $X_i \leftarrow \text{Explore}(X', P)$  // Create new solution set
5    $\Delta Y_i \leftarrow \text{Model.Predict}(X_i)$ 
6   // Select best  $P$  solution from current union previous.  $\Delta Y'$ ,
    $X' \leftarrow \text{Select}(X' \cup X_i, \Delta Y_i \cup \Delta Y', P)$ 
7 end
8 // Select the top  $k$  solution among  $\Delta Y'$ 
9  $(X', \Delta Y') \leftarrow \text{Top}(X', \Delta Y', k, F_{obj})$ 

```

### 5.1 Click Prediction Performance

In the click prediction experiments, we used 10 million sample of data that spans a week time period from Bing PC traffic of January 2018. For the experimental setup, the size of training data is gradually increased from 100K to 3.2M. In each data point, the subset of data excluding the training set is used for testing.

Figure 5 presents the cumulative error for click prediction models with respect to Equation 13 where  $N$  is the number of test impressions.

$$E_{\text{cumulative}} = \frac{|\sum_{i=1}^N y_i - \sum_{i=1}^N \hat{y}_i|}{\sum_{i=1}^N y_i} \quad (13)$$

For the cumulative error, we observed that Tree Boosting is slightly better than Bayesian Probit when the data size is smaller than 3M samples. Another observation is that when data size becomes larger than 3M samples, the error rate is getting stable and converges for homogenous traffic like Bing PC. This behavior is very similar to other homogenous traffic slices that are tuned by Genie.

While the cumulative error gives general idea for click prediction quality on the aggregated level, it does not reflect the success of individual impression predictions. The normalized Log Loss error in Equation 14 gives better estimation of impression level prediction accuracy. In the second experiment, we compared the Log Loss error of Bayesian Probit and Tree Boosting (Figure 6). We observe that Tree Boosting is slightly better than Bayesian Probit in terms of Log Loss Error. The convergence behavior is very similar to cumulative prediction error after 3M samples.

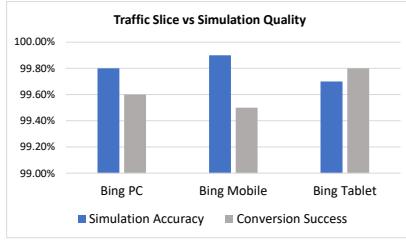


Figure 8: Simulation Quality

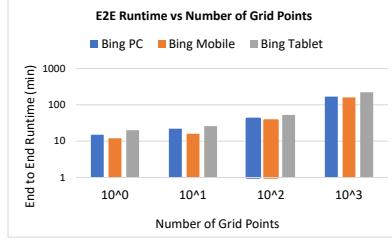


Figure 9: Simulation Runtime

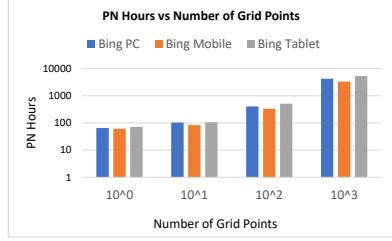


Figure 10: Simulation PN Hours

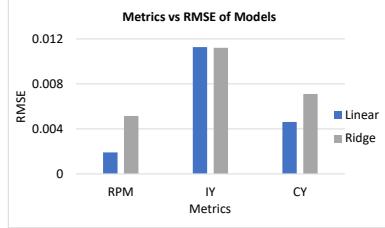


Figure 11: Model Accuracy

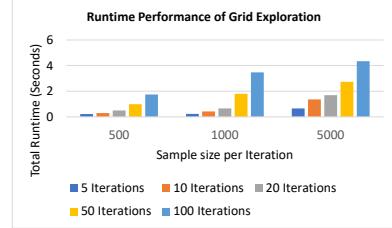


Figure 12: Grid Exploration Performance

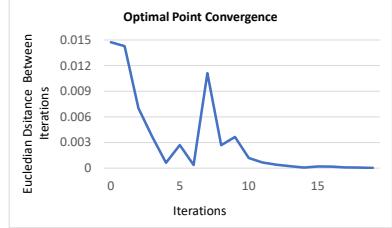


Figure 13: Optimal Point Convergence

$$E_{logloss} = \frac{1}{N} \sum_{i=1}^{i=N} \psi(y_i, \hat{y}_i) \quad (14)$$

In the last experiment, we compare the runtime performance of the model training part of Bayesian Probit and Tree Boosting. The training part for the Bayesian Probit runs on cosmos and integrated into Auction Simulation Job. The training job for Tree Boosting is running in a single box on Microsoft Internal Workflow Platform outside the Auction Simulation Job. Figure 7 presents the runtime comparison of training jobs for these two models. We observe that the training part of both model is very fast and have similar performance. However, the current Tree Boosting implementation is not compatible with Cosmos and running outside the cosmos data cluster. Therefore, the significant amount of time (10 min - 2 hours) is spent on transferring training data between data centers (Figure 7 excludes this data). Because of this, Tree Boosting is only preferred for cases when significant feature engineering is involved (non continuous features) or noisy traffic slices in terms of user behavior and feature space (traffic types where Bing is not publisher).

## 5.2 Auction Simulation Experiments

In this part, we analyze the runtime performance and replay quality of the Auction Simulation step. For the experimental purpose, we use the 3 different data samples collected from Bing PC, Bing Mobile and Bing Tablet traffic slices in January 2018.

In this experiment, we prepare 12 job configs for 3 traffic slices with 4 different grid settings. The performance results only reflects simulation phase of the auction simulation job that is executed on Cosmos platform. The Auction Simulation component runs on 2000 processing node up to certain number of data samples cross the number of input grid points. Once the number of input grid points and data size increase, we gradually increase the number of processing nodes based on heuristics which is calibrated periodically. In the current experiment, we used 4 different grid size on exponential scale from  $10^0$  to  $10^3$ . The first 3 data points correspond to

simulation run with 2000 nodes and the grid size of  $10^3$  corresponds to simulation run with 10000 processing nodes.

Figure 8 presents the simulation quality for this experiment. Simulation accuracy represents the percentage of requests that can be replayed correctly in offline among converted logs. Conversion success represents the percentage of requests which are logged properly and can be converted into the data structure that online library expects. As Figure 8 suggests that, both simulation accuracy and conversion success are constantly above 99% which is needed for accurate counterfactual estimation. Figure 9 and Figure 10 present runtime performance metrics. As the number of grid points increases, the runtime and PN hours increase too. Another observation is that the tangent of end to end time is not sharp as in total PN hours till some point ( $10^3$ ). This implies that as long as there is an available resource for parallelization, the end to end runtime is less impacted compared to the total processing time.

## 5.3 Grid Exploration

For the grid exploration, we use one of the specific tuning scenario for Bing PC slice with 2K grid points that spans a week of data in Jan 2018. Through this section, Grid Exploration experiments run in a single machine that has Intel i7 CPU with 4 cores and 32 GB RAM.

In the first experiment, we compared the prediction accuracy of Linear and Ridge regression. For both models, we tried 3 different versions that use linear, quadratic and cubical polynomial feature space. Here, we don't present higher polynomial feature space than cubical since increasing the order of polynomial feature space more causes over fitting and does not contribute prediction accuracy significantly. For reporting purpose, we only compare the most common metrics that are used for tuning decisions: Revenue per mille (RPM), click yield (CY) and impression yield (IY). Figure 11 presents the Root Mean Square Error (RMSE) comparison of Ridge and Linear regression. In this experiment, we use 3-fold cross validation where we randomly divide data into 3 subset and

2 subset is used for training and rest is used for prediction. Our first finding is that Linear regression model yields better prediction performance compared to Ridge regression. We also discovered that cubical features yields best performance in majority of the tuning scenarios.

In the second and third experiments, we analyze the performance of the optimization algorithm mentioned in Section 4.2. Figure 12 presents the runtime performance with respect to the number of solutions explored in each step and the number of maximum iterations. As we increase the number of iterations and solution space, the runtime increases linearly. In terms of number of solutions, increasing the solution space in each iteration does not contribute significantly after 5000 solutions. Another interesting results is that, the grid exploration method with 5000 solution space converges very quickly after 20 iterations (Figure 13). This is indeed a very common behavior that we observed in thousands of tuning jobs where grid size coming from Genie jobs are less than 5K.

#### 5.4 Comparison of Policy Estimators

In this section, we compare counterfactual estimation performance of Genie with Importance Sampling based estimator [2] which is one of the most common observational approaches that is utilizing inverse propensity weight estimation. Importance Sampling computes any metric  $y$  based on given online distribution of system parameters  $P(x)$  over sample of  $N$  data points as follows:

$$Y = \int_x yP(x) \approx \frac{1}{N} \sum_{i=1}^N y_i \quad (15)$$

For the counterfactual distribution  $P^*(x)$  of online parameters, the Counterfactual KPI  $Y^*$  can be estimated as follows where the  $w(x_i)$  corresponds to ratio of counterfactual distribution over proposal distribution for current sample  $x_i$ :

$$Y^* = \int_x y \frac{P^*(x)}{P(x)} P(x) \approx \frac{1}{N} \sum_{i=1}^N \frac{P^*(x_i)}{P(x_i)} y_i = \frac{1}{N} \sum_{i=1}^N w(x_i) y_i \quad (16)$$

Since Importance Sampling needs online randomization of tunable parameter set  $x$  under certain distribution of  $P(x)$ , we choose less risky tuning scenarios that is used for periodic tuning of existing model parameters. For comparison purpose, we choose 5 tuning interval for the same problem between April and May 2018 for Bing PC traffic slice. We also created randomized experiment with Gaussian distribution on certain set of parameters that are used for controlling the same model. After that, we tuned same parameter set periodically 5 times via importance sampling on randomized logs with a certain objective function and a set of k proposal distributions  $PD = \{P_1^*(x), P_2^*(x), \dots, P_k^*(x)\}$ . At the end of each tuning period, the best solution from corresponding tuning jobs is applied to treatment traffic under  $A/B$  experiment setup for certain amount of testing time. For comparison purpose, we also estimated the KPI impact of selected solutions (selected proposal distribution) for each time interval by using Genie tunings on exactly same data.

Table 1 presents comparison of Genie Estimator with Importance Sampling based Estimator. The first and second rows represent the average error from tunings based on data collected from randomized experiment during tuning interval. The error in this context is the

**Table 1: Comparison with Importance Sampling**

Method	RPM	MLIY	CY	CPC
IS (Historical)	1.27%	0.41%	0.39%	1.14%
Genie (Historical)	<b>1.16%</b>	<b>0.32%</b>	<b>0.37%</b>	<b>0.93%</b>
IS (Regression)	0.90%	0.36%	<b>0.24%</b>	0.98%
Genie (Regression)	<b>0.88%</b>	<b>0.25%</b>	0.27%	<b>0.66%</b>

distance between predicted KPI delta of  $k^{th}$  interval and actual KPI delta of  $(k+1)^{th}$  interval that is computed from the  $A/B$  experiment. The last two rows represent the average error from regression tunings where we apply same counterfactual on control traffic of  $A/B$  tests on each interval (predicted and actual KPIs belong to same interval). Our first observation is that Genie outperforms Importance Sampling estimator on the space (MLIY) and pricing metrics (CPC) on noisy traffic slices since the system layer in causal graph has significant impact on the outcome of these variables. Another interesting observation is that both estimator performs closer for Click Yield and RPM the outcome of which is significantly impacted by user behaviors directly or indirectly. This behavior is indeed pretty similar in many other tunings scenarios where we observe that Genie is more successful on predicting metrics that has more noise due to system related changes like frequent policy updates. Our final observation is that both tunings give better result if they use the same control data for tuning with  $A/B$  testing time period which are listed as regression tunings in Table 1.

#### 5.5 Tuning New Policies with Genie

In this section, we present experimental results on tuning new policies with Genie for Bing PC traffic. Tuning new policies is the most risky tuning scenarios since there may not be enough knowledge on hyperparameter space. Thus, creating  $A/B$  testing or randomized experiment could be very costly due to cold start problem.

For this particular experiment, we used periodic combo tuning scenario where parameters that control new models from multiple teams are tuned jointly to achieve global optimization goals. To find the best operating point for this tuning task, Genie user submitted more than 6 jobs. After tunings, user published the best operating point and created two online experiment for  $A/B$  testing which are randomly sampled versions of the current production Bing PC traffic slice. Then, the first experimental slice is updated to the new operating point by applying optimal parameter settings as a treatment and the second clone is kept as a control slice for  $A/B$  testing comparison.

For KPI validation purpose, we also submit Genie job on control traffic of  $A/B$  test by applying same treatment in simulation for comparison purpose. The results of current experiments are given in Table 2 for RPM (Revenue Per Mille), CY (Click Yield), MLIY (Main Line Impression Yield). Our first observation is that Genie results are consistent with real traffic and metrics of interests are directionally correct. Our second observation is that results are getting better when data for real traffic and logs of Genie job are similar and sampled from same dates (row 1 and row 2 in Table 2). When we shift temporal dimension to the past, we started to observe more

**Table 2: End to End Tuning Experiments**

Job	Dates	$\Delta RPM$	$\Delta CY$	$\Delta MLIY$
Real Delta	01/31-02/02	+1.60%	+0.14%	+4.74%
Genie Tuning	01/31-02/02	+1.73%	+0.07%	+5.13%
Genie Tuning	01/24-01/30	+1.81%	+0.04%	+5.33%

discrepancy (delta of row 1 and row 3) since user and advertiser behaviors are getting different as training data gets older compared to *A/B* testing time frame.

## 6 CONCLUSION REMARKS

In this paper, we propose an offline counterfactual policy estimation framework called Genie to optimize the Sponsored Search marketplace. From the experimental results on Bing traffic, we show that Genie can be used to tune completely new policy changes or parameters of existing policies for noisy traffic slices. We also discovered that Genie can yield better results compared to existing observational approaches for metrics the outcome of which depend more on policy changes. For metrics the outcome of which depends on user behavior, Genie can also produce similar KPI results compared to observational approaches like Importance Sampling. Thus, system simulation + user modeling mechanism that Genie is utilizing could be applied to similar problem contexts where training data has bias due to frequent policy updates.

## ACKNOWLEDGMENTS

The authors would like thank Emre Kiciman for feedbacks on improving the content of this paper, Tianji Wang and Qian Yao for their help on running Policy Estimators comparison experiments. Finally, authors would like to thank Rukmini Iyer, Eren Manavoglu, Patrick Jordan, Manish Agrawal, Patryk Zaryjewski, Anton Schwaighofer, Sendil Arunachalam, Pei Jiang, Suyesh Tiwari, Lohith Shesharam, Ranjita Naik, Marcin Machura, Urur Dogan, Elon Portugaly, Jitendra Ajmera, Aniruddha Gupta, Dilan Gorur, Debabrata Sengupta, Eugene Larchyk, Tommy Tan, Xiaoliang Ling, Thomas Borchert, many talented scientists and engineers in Microsoft Bing Ads Team for their help on Genie implementation and feedbacks for many features of Genie.

The auction described in the paper is for illustration only and it is not necessarily the one that Bing uses.

## REFERENCES

- [1] Susan Athey. 2015. Machine learning and causal inference for policy evaluation. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. ACM, 5–6.
- [2] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. *The Journal of Machine Learning Research* 14, 1 (2013), 3207–3260.
- [3] Andrei Broder, Peter Ciccolo, Evgeniy Gabrilovich, Vanja Josifovski, Donald Metzler, Lance Riedel, and Jeffrey Yuan. 2009. Online expansion of rare queries for sponsored search. In *Proceedings of the 18th international conference on World wide web*. ACM, 511–520.
- [4] Ronnie Chaiken, Bob Jenkins, Per-Åke Larson, Bill Ramsey, Darren Shakib, Simon Weaver, and Jingren Zhou. 2008. SCOPE: easy and efficient parallel processing of massive data sets. *Proceedings of the VLDB Endowment* 1, 2 (2008), 1265–1276.
- [5] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 785–794.
- [6] Daniel C Fain and Jan O Pedersen. 2006. Sponsored search: A brief history. *Bulletin of the Association for Information Science and Technology* 32, 2 (2006), 12–13.
- [7] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [8] Peter W Glynn and Donald L Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35, 11 (1989), 1367–1392.
- [9] Gagan Goel, Vahab Mirrokni, and Renato Paes Leme. 2015. Polyhedral clinching auctions and the adwords polytope. *Journal of the ACM (JACM)* 62, 3 (2015), 18.
- [10] Avi Goldfarb. 2014. What is different about online advertising? *Review of Industrial Organization* 44, 2 (2014), 115–129.
- [11] Thore Graepel, Joaquin Q Candela, Thomas Borchert, and Ralf Herbrich. 2010. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th international conference on machine learning (ICML-10)*. 13–20.
- [12] Keisuke Hirano and Guido W Imbens. 2001. Estimation of causal effects using propensity score weighting: An application to data on right heart catheterization. *Health Services and Outcomes research methodology* 2, 3-4 (2001), 259–278.
- [13] Ron Kohavi, Alex Deng, Brian Frasca, Toby Walker, Ya Xu, and Nils Pohlmann. 2013. Online controlled experiments at large scale. In *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11–14, 2013*. 1168–1176. <https://doi.org/10.1145/2487575.2488217>
- [14] Ron Kohavi, Randal M Henne, and Dan Sommerfield. 2007. Practical guide to controlled experiments on the web: listen to your customers not to the hippo. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 959–967.
- [15] Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M Henne. 2009. Controlled experiments on the web: survey and practical guide. *Data mining and knowledge discovery* 18, 1 (2009), 140–181.
- [16] Lihong Li, Shunbao Chen, Jim Kleban, and Ankur Gupta. 2015. Counterfactual estimation and optimization of click metrics in search engines: A case study. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, 929–934.
- [17] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, et al. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1222–1230.
- [18] Matthew Richardson, Ewa Dominowska, and Robert Ragno. 2007. Predicting clicks: estimating the click-through rate for new ads. In *Proceedings of the 16th international conference on World Wide Web*. ACM, 521–530.
- [19] Nir Rosenfeld, Yishay Mansour, and Elad Yom-Tov. 2017. Predicting Counterfactuals from Large Historical Data and Small Randomized Trials. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*. 602–609. <https://doi.org/10.1145/3041021.3054190>
- [20] Amin Sayedi, Kinshuk Jerath, and Kannan Srinivasan. 2014. Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising. *Marketing Science* 33, 4 (2014), 586–608.
- [21] Tobias Schnabel, Paul N. Bennett, Susan T. Dumais, and Thorsten Joachims. 2018. Short-Term Satisfaction and Long-Term Coverage: Understanding How Users Tolerate Algorithmic Exploration. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, February 5-9, 2018*. 513–521. <https://doi.org/10.1145/3159652.3159700>
- [22] Adith Swaminathan and Thorsten Joachims. 2015. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research* 16 (2015), 1731–1755. <http://dl.acm.org/citation.cfm?id=2886805>
- [23] Adith Swaminathan and Thorsten Joachims. 2015. The self-normalized estimator for counterfactual learning. In *Advances in Neural Information Processing Systems*. 3231–3239.
- [24] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.
- [25] Hal R Varian and Christopher Harris. 2014. The VCG auction in theory and practice. *American Economic Review* 104, 5 (2014), 442–45.
- [26] Wanrong Xu, Eren Manavoglu, and Erick Cantú-Paz. 2010. Temporal click model for sponsored search. In *Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2010, Geneva, Switzerland, July 19–23, 2010*. 106–113. <https://doi.org/10.1145/1835449.1835470>
- [27] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1077–1086.