Automated Tuning of Ad Auctions

Dilan Gorur, Debabrata Sengupta, Levi Boyles, Patrick Jordan, Eren Manavoglu,

Elon Portugaly, Meng Wei, Yaojia Zhu
Microsoft
{dgorur, desengup, leboyles, pajordan, ermana, elonp, mewe, yaozhu}@microsoft.com

Abstract

We present an application of counterfactual reasoning for completely automated tuning of online ad auctions. We discuss practical considerations for applying the approach on a major search engine and present results showing parity between a reliable simulation engine and our approach that is orders of magnitude more efficient computationally.

1 Online Ad Auctions

Selection, pricing, and allocation of ads on search result pages is determined via an auction that runs in real time whenever a query is submitted to a search engine. In each instance, relevant ads are selected to participate and allocation and pricing rules are subsequently executed to determine the placement and click-contingent price for each ad. Varian [3] popularized an abstraction of this auction known as Generalized Second Price, whereby ads are allocated in descending order of rank score and priced such that the payment per click is equal to the minimum bid required to maintain its current position according to the rank score function. Varian describes the rank score of ad i with bid b_i as $S_i(b_i) = q_i b_i$, where q_i is some advertiser dependent quality score. In this paper, we consider a more generalized form of the rank score function, parameterized by auxiliary parameters $\{\lambda_j\}_1^K$, that influence pricing and allocation via advertiser, user and query dependent quality scores $\{q_j\}_1^K$, namely $S_i(b_i|\lambda, \mathbf{q})$.

The auxiliary parameters control the trade-off between different key performance indicators (KPIs) such as revenue (RPM), click through rate (CTR), quick back rate (QBR), and ad coverage. The parameter values can be adjusted to achieve a particular operating point. Changes in the models that feed into the auction as well as changes in the marketplace make it necessary for the parameters to be readjusted, a process we refer to as *tuning*. Typically, tens of experiments need to be regularly tuned per market per device per month with many markets worldwide. The number of tunings required per month and the dynamic nature of marketplace requirements makes this a daunting task. This work builds upon [1] and focuses on the practical aspects of utilizing counterfactual reasoning in an end-to-end automated ad auction tuning pipeline for a major search engine. The idea is to use a distribution over the parameters at serve time, which enables asking "what-if" questions and predicting what the expected KPIs would be given alternative marketplace parameter settings. Here, we focus on the practical aspects of this approach and refer the reader to [1] for details on the basic technical details.

2 Quantifying and Controlling Randomization Cost

We make use of randomized auction parameters which are then used in conjunction with importance sampling in order to perform counterfactual estimation. For simplicity, we restrict our attention to the case with proposal $\mathcal{N}(\mu, \Sigma)$. Choice of the proposal distribution p for which to sample the

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

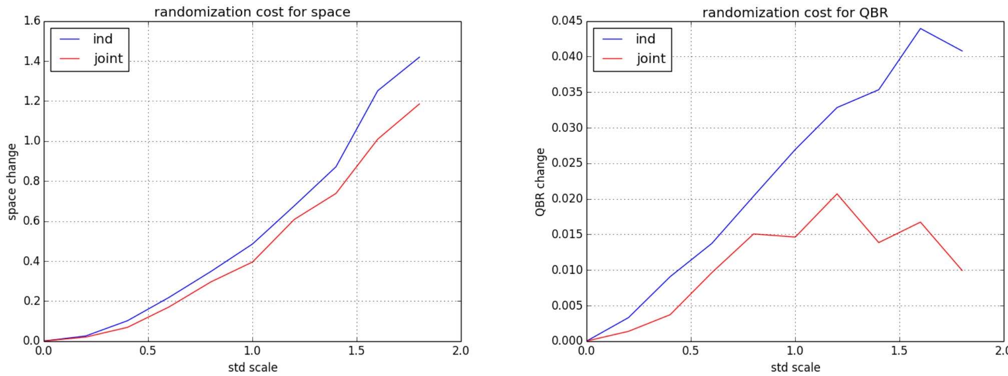

Figure 1: Quantifying randomization cost using Monte Carlo simulations. The choice of the proposal distribution should be guided by the desired amount of exploration and the amount of deviation from the non-randomized experiment's KPIs that can be afforded. Randomization cost can be reduced using parameter correlations.

auction parameters from affects the exploration power of the method. Proposals with larger variance will allow a larger range of counterfactuals to be considered. However, randomizing the parameters associated with a real-time auction has an impact on the overall KPIs as well, so there is a trade-off that must be considered. A larger variance will result in parameter samples further away from the current operating point. If the KPIs are not linear functions of the parameters, then the KPIs will be impacted in some way with the effect of the impact (typically) increasing with the size of the variance. The correlation structure of the parameters will impact the KPIs as well. Specifically, it is natural to correlate parameters that impact KPIs similarly so that more samples are exploring interesting regions of the counterfactual space. The corresponding joint covariance can have a different impact on the KPIs than that of an independent covariance with the same marginals as the joint.

We utilized an existing auction replay and user simulation platform to obtain *simulated counterfactuals* in order to compute the correlation structure of the marketplace parameters, estimate the impact of randomization on the KPIs and finally decide on a covariance for online experiments. For computing the correlation structure, we simulated counterfactual KPIs for a large grid of marketplace parameters. We applied rather lose constraints on the counterfactual KPIs to eliminate points that lead to very poor operating points. We treated the parameter settings for the remaining points as possibly viable for a tuning scenario, and computed the correlation on these points. Next, we tested the impact of randomization using independent and correlated distributions with increasing variance.

Let f be our KPI of interest and x our auction parameters. If the marketplace parameter distribution is $p(x)$ and the distribution of the KPI given a parameter setting is $p(f|x)$, then we are interested in the quantity

$$E_{p(f)}[f] = E_{p(f,x)}[f] = \int fp(x)p(f|x)dxdf = \int dxp(x) \int fp(f|x)df = E_{p(x)}[E_{p(f|x)}[f|x]]. \tag{1}$$

The simulation platform performs its estimations given user queries by training user click models and applying these to simulated auctions and clicks with computationally intensive simulation models, providing a reliable estimate of the KPI: $\hat{f}(x) \approx E_{p(f|x)}[f|x]$. Given this approximation, we can construct a Monte Carlo estimate of the KPI with

$$E_{p(f)}[f] \approx E_{p(x)}[\hat{f}(x)] \approx \sum_{x_i \sim p(x)} \hat{f}(x_i). \tag{2}$$

We estimated the correlation structure Σ of the parameters using simulation data as explained above. For Σ_1, a diagonal covariance matrix, and Σ_2, a full covariance matrix, such that Σ_1 and Σ_2 have the same marginals, we evaluated (2) for $p_1 = \mathcal{N}(\mu, c^2\Sigma_1)$ and $p_2 = \mathcal{N}(\mu, c^2\Sigma_2)$ for a range of values for c ($c = 0$ corresponds to no randomization). Figure 1 compares the cost of randomization for different values of c, for the two matrices, Σ_1 (left) and Σ_2 (right). We found that, as expected, larger values of c correspond to larger impact on the KPIs. Additionally, we see that jointly randomizing the auction parameters reduces the cost of randomization.

Table 1: Randomization Impact Comparison for Online Experiments(percent change)

	RPM	QBR	Space	Impression
Ind. $\gamma = 1$	0.38	-0.15	0.48	0.51
Ind. $\gamma = 0.5$	0.21	0.08	0.22	0.28
Joint $\gamma = 1$	-0.15	0.21	0.09	0.11

In order to have a tighter grip on the randomization cost, we also introduce a randomization rate parameter that can be used to reduce the impact of randomization on the KPIs. To this end, we specify the flight distribution as $p = (1 - \gamma)\delta_\mu + \gamma N(\mu, \Sigma)$, so that the impact of randomization can be reduced by reducing γ. In this way, we can reduce the impact of randomization without reducing the size of our effective counterfactual search space (as would be the case with reducing c), at the cost of requiring more time to gather the same amount of randomized data.

Finally, we ran online experiments and compared their KPIs to an unrandomized baseline (with all means and centers the same). See Table 1. This experiment verifies that joint randomization indeed has lower impact than independent randomization. Some KPIs are closer to baseline even compared to independent randomization run with $\gamma = 0.5$. Thus, we conclude that utilizing correlation structure of the parameters leads to a more efficient exploration.

3 Estimation Uncertainty

We use importance sampling to estimate counterfactual KPIs using a weighting of the KPIs observed using a particular sampling distribution in the online serving environment. Confidence intervals computed using importance sampling may be underestimated when the sample is not representative of the population. We use effective sample size as a diagnostic to tell if the importance sampling estimate is reliable. Noting that revenue or click based KPIs are generally larger variance than impression based KPIs, we compute both overall effective sample size and KPI-based effective sample size using

$$\text{ESS}(f) = \frac{[\sum_i w_i f(x_i)]^2}{\sum_i w_i^2 f(x_i)^2}, \tag{3}$$

and consider counterfactuals that have ESS lower than a threshold to be unreliable.

4 Increasing Computational Efficiency

When both the proposal distribution and counterfactual distributions are Normally distributed with the same covariance, we can compute the ESS that would occur asymptotically for each counterfactual and pre-filter them before performing importance sampling.

Consider samples $x \sim q = N(0, \Sigma_x)$ and a counterfactual distribution $p = N(\mu_x, \Sigma_x)$, giving weights $w(x) = \frac{p(x)}{q(x)}$. Then we have

$$\log w(x) = x^T \Sigma_x^{-1} \mu_x - \frac{\mu_x^T \Sigma_x^{-1} \mu_x}{2}. \tag{4}$$

As the above is a linear map on x and $x \sim N(0, \Sigma_x)$, we have

$$\log w(x) \sim N(-\frac{\mu_x^T \Sigma_x^{-1} \mu_x}{2}, \mu_x^T \Sigma_x^{-1} \mu_x). \tag{5}$$

Thus the weights are Log-Normal with mean parameter $\mu = -\frac{\mu_x^T \Sigma_x^{-1} \mu_x}{2}$ and variance parameter $\sigma^2 = \mu_x^T \Sigma_x^{-1} \mu_x$. As

$$E[w] = \exp\left(\mu + \frac{\sigma^2}{2}\right) = \exp(0) = 1,$$
$$E[w^2] = \exp\left(2\mu + 2\sigma^2\right) = \exp\left(\mu_x^T \Sigma_x^{-1} \mu_x\right).$$

Thus, asymptotically we expect the ESS on N samples to converge to

$$ESS \to N e^{-\mu_x^T \Sigma_x^{-1} \mu_x}. \tag{6}$$

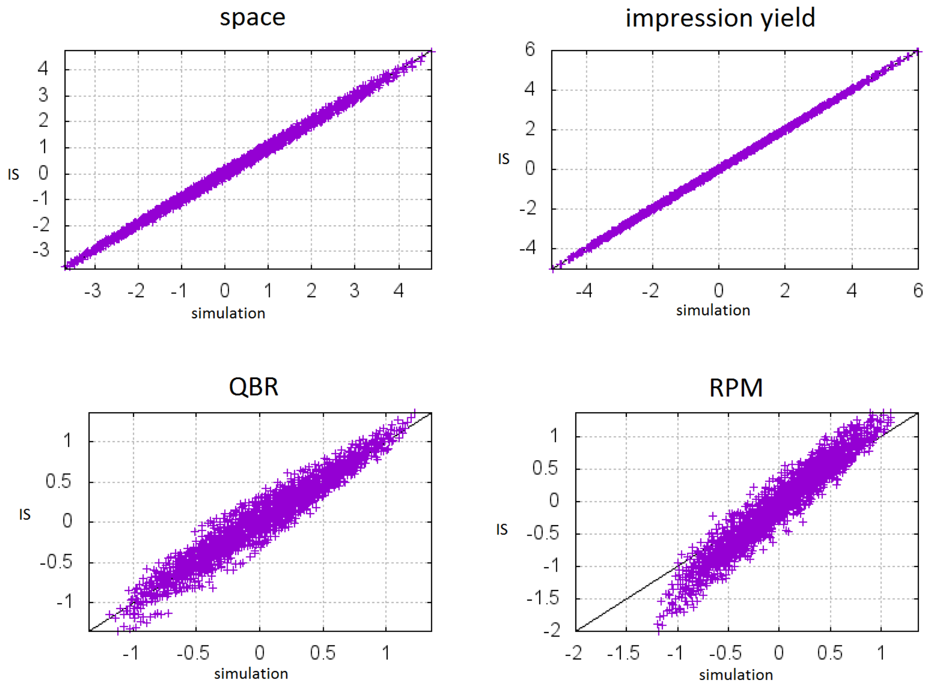

Figure 2: Comparing counterfactual estimates using IS to the simulations obtained from auction replay and user simulation platform.

It is interesting to note that ESS depends on the shift in the mean, not on the dimension of the distribution. Thus, we can reliably evaluate the counterfactual KPIs for small changes to a large set of parameters as well as larger changes to a small set of parameters. We use this observation and asymptotic ESS values to guide choosing the search space and avoid computations on parameter settings that would not return reliable estimates.

5 Conclusion

We have presented practical aspects of applying an efficient counterfactual estimation method utilizing importance sampling. It confers many advantages over other counterfactual estimation methods (simplicity of implementation, faster throughput, unbiasedness) at the cost of introducing some randomness into the real time auction. Comparison to the auction replay and user simulation platform in Figure 2 shows that we produce similar estimates at a fraction of the computational cost.

Using this approach, we can explore only a limited counterfactual space. However, the limited search space is not such a burden as tunings can be performed more frequently, even daily, in order to keep up with changing marketplace demands. Combined with a suitable KPI constrained optimization policy, we have run automated daily tunings. Early results show that we are able to successfully tune and achieve marketplace target KPIs. Such an automated workflow "removes a human from the loop," improving the health and reliability of the tuning system, as it is less prone to human error.

References

[1] L. Bottou, J. Peters, J. Quiñonero Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(Nov):3207–3260, 2013.

[2] A. Swaminathan and T. Joachims. The self-normalized estimator for counterfactual learning. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3231–3239. Curran Associates, Inc., 2015.

[3] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2006.