Evaluating User Targeting Policies: Simulation Based on Randomized Experiment Data

Joel Barajas*
UC, Santa Cruz
Santa Cruz CA, USA
jbarajas@soe.ucsc.edu

Ram Akella†
School of Information
UC, Berkeley CA, USA
akella@
ischool.berkeley.edu

Marius Holtan
AOL Research
Palo Alto CA, USA
marius.holtan@teamaol.com

ABSTRACT

We propose a user targeting simulator for online display advertising. Based on the response of 37 million visiting users (targeted and non-targeted) and their features, we simulate different user targeting policies. We provide evidence that the standard conversion optimization policy shows similar effectiveness to that of a random targeting, and significantly inferior to other causally optimized targeting policies.

Categories and Subject Descriptors

J.1 [**Computer Applications**]: Administrative Data Processing—*Marketing*; G.3 [**Mathematics of Computing**]: Probability and Statistics—*Experimental Design*

Keywords

Causal Attribution; Targeted Advertising; A/B Testing

1. INTRODUCTION

The use of randomized experiments is becoming the standard practice to accurately measure the ad casual effect on user conversions [2]. Given that randomized experiments are expensive, the generated data should be leveraged as much as possible. However, the use of this data has been limited to the ad effectiveness estimation only. On the other hand, user targeting development has focused largely on optimizing user conversions by serving ads to the users who are more likely to convert [3]. Often the evaluation of these algorithms is based on the prediction power of conversions, which are likely to be not caused by the campaign [2]. This practice often leads to large discrepancies when these algorithms are tested in a randomized experiment.

We propose a simulator that leverages the data of randomized experiments by considering all the visiting users to the publisher websites [1]. We fit the user conversion response of the campaign/placebo ad exposures (targeted users), and the response of those who are not targeted. Based on the data of a randomized experiment for 37 million users, 8 million targeted users, and user demographic features, we simulate the standard conversion optimization policy and three targeting algorithms based on the ad average causal effect.

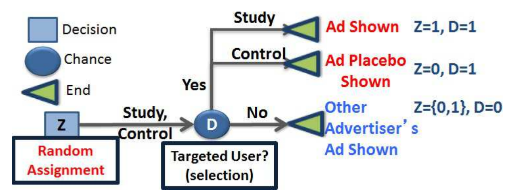

Figure 1: Randomized design for all visiting users.

2. METHODOLOGY

The standard practice to estimate the ad causal effect is to run a randomized experiment where the online visiting users are randomly assigned to the control or the study treatment arms. For those assigned to the study group, the campaign ad is displayed, while a placebo ad is displayed to the users of the control group. In practice, a placebo campaign, which replicates the targeting performed by the focal campaign, is run to display the placebo ads. Fig. 1 depicts this process.

We define the following indicator variables for each user i: Z_i for control/study assignments $\{0,1\}$; D_i for non-targeted/targeted users $\{0,1\}$; Y_i for non-converting/converting users $\{0,1\}$; and X_i for feature segments defined to be finite and countable. We find the user counts by segments, $N^y_{dz}|X_i$ given $D_i = d$, $Z_i = z$, $Y_i = y$, X_i, leading to the set:

$$N^{obs} = \{N^y_{dz}|X_i : \forall d \in \{0,1\}, \forall z \in \{0,1\}, \forall y \in \{0,1\}, \forall X_i\}$$

The cardinality of N^{obs} becomes: $\#\{N^{obs}\} = 8 \times \#\{X_i \forall i\}$.

We use the data, N^{obs}, to simulate a given targeting function, $F_{targ}(X_i)$, based on Algorithm 1. We model the user response to the campaign and the placebo ads, $P(Y_i = 1|D_i = 1, Z_i = z, X_i) = \theta_{1z}|X_i : \forall z \in \{0,1\}$, as well as the response of the non-targeted population, $P(Y_i = 1|D_i = 0, Z_i = z, X_i) = \theta_{0z}|X_i : \forall z \in \{0,1\}$, through a probit transformation as illustrated by steps 3–4 of Algorithm 1. Here, glmfit($[N^1|X, N^0|X]$) represents standard probit regression fitting given the vectors of successes and failures $N^1|X_i, N^0|X_i$, and feature vector X_i. We consider the observed targeted users as a fixed campaign budget (N^{budget}_{1z}, step 6). This budget is consumed by the user targeting of step 12, which includes the probability of user segments ($P(X_i)$). The *min* function enforces the visiting population segment constraints ($N^{Visit}_{remain}|X_i$). The while loop of steps

*Main contact.
†Also: UC, Santa Cruz, akella@soe.ucsc.edu

WWW 2015 Companion, May 18–22, 2015, Florence, Italy.
ACM 978-1-4503-3473-0/15/05.
http://dx.doi.org/10.1145/2740908.2742778.

Algorithm 1 User Targeting Simulation

1: **Input:** Targeting function $F_{targ}(X_i)$, User Counts $N_z^{obs} = \{N_{dz}^y|X_i : \forall d \in \{0,1\}, \forall y \in \{0,1\}, \forall X_i\}$.
2: **Output:** Aggregated User Counts After Targeting $N_{z,agg}^{new} = \{N_{dz}^{y,new} : \forall d \in \{0,1\}, \forall y \in \{0,1\}\}$
3: Set $[\hat{\gamma}_{0z}, \hat{\gamma}_{1z}]$=[glmfit$\left([N_{0z}^1|X_i, N_{0z}^0|X_i], \forall X_i\right)$, glmfit$\left([N_{1z}^1|X_i, N_{1z}^0|X_i], \forall X_i\right)$] // *Probit Approximation*
4: Set $[\hat{\theta}_{0z}, \hat{\theta}_{1z}]|X_i = [\Phi(X_i'\hat{\gamma}_{0z}), \Phi(X_i'\hat{\gamma}_{1z})], \quad \forall X_i$ // *Observed Conversion Propensity*
5: Set $N_z^{Visit}|X_i = N_{1z}^1 + N_{1z}^0 + N_{0z}^1 + N_{0z}^0|X_i, \quad \forall X_i$ // *Audience per Segment X_i*
6: Set $N_{1z}^{budget} = \sum_{\forall X_i}(N_{1z}^1 + N_{1z}^0)|X_i$ // *Observed Budget*
7: Set $N_{1z}^{1,new}|X_i = N_{1z}^{0,new}|X_i = 0, \quad \forall X_i$ // *Set Counts*
8: Set $N_{remain}^{budget} = N_{1z}^{budget}$ // *Initialize Remaining Budget*
9: **while** $N_{remain}^{budget} > 0$ **do**
10: Set $P(X_i) = N_{remain}^{Visit}|X_i / \sum_{\forall X_i} N_{remain}^{Visit}|X_i, \quad \forall X_i$
11: Set $\lambda = N_{remain}^{budget} / \left(\sum_{\forall X_i} N_{remain}^{budget} \times F_{targ}(X_i) \times P(X_i)|X_i\right)$ // *Budget Multiplier*
12: Set $\left[N_{1z}^{1,new}, N_{1z}^{0,new}\right]|X_i = \left[N_{1z}^{1,new}, N_{1z}^{0,new}\right]|X_i + \min\left(\lambda \times F_{targ}(X_i) \times N_{remain}^{budget} \times P(X_i), N_{remain}^{Visit}|X_i\right) \times \left[\hat{\theta}_{1z}, 1 - \hat{\theta}_{1z}\right]|X_i, \quad \forall X_i$ // *Target Users*
13: Set $N_{remain}^{Visit}|X_i = N_z^{Visit} - (N_{1z}^{1,new} + N_{1z}^{0,new})|X_i, \quad \forall X_i$ // *Remaining Audience*
14: Set $N_{remain}^{budget} = N_{1z}^{budget} - \left(\sum_{\forall X_i}[N_{1z}^{1,new} + N_{1z}^{0,new}|X_i]\right)$ // *Remaining Budget*
15: **end while**
16: Set $\left[N_{0z}^{1,new}, N_{0z}^{0,new}\right]|X_i = N_{remain}^{Visit} \times \left[\hat{\theta}_{0z}, 1 - \hat{\theta}_{0z}\right]|X_i, \quad \forall X_i$ // *Non-Targeted User Counts*
17: Set $N_{z,agg}^{new} = \left\{\sum_{\forall X_i} N_{dz}^{y,new}|X_i : \forall d \in \{0,1\}, \forall y \in \{0,1\}\right\}$ // *Aggregate User Counts*

9–15 re-distributes the remaining budget in case $N_{remain}^{Visit}|X_i$ is exhausted for any segment. We aggregate the user counts over X_i to generate the four counts given $Z_i = z$: $N_{z,agg}^{new} = \{N_{dz}^{y,new} : \forall d \in \{0,1\}, \forall y \in \{0,1\}\}$. This simulation is run for both treatment arms $z \in \{0,1\}$ independently, and the ad effect is measured based on a t-test of ATE=$E(Y_i|D_i = 1, Z_i = 1) - E(Y_i|D_i = 1, Z_i = 0) = \theta_{11} - \theta_{10}$, using $\{N_{1z}^{y,new} : \forall z \in \{0,1\}, \forall y \in \{0,1\}\}$.

3. RESULTS

We consider the user features: age, gender and income; segmented by value ranges (finite and countable). The campaign running time is two weeks. For $Z_i = 1$, the total and targeted population sizes are 18.74 and 4.01 million. For $Z_i = 0$, the total and targeted population sizes are 18.70 and 4.09 million. We consider the missing values as a feature value (81.4% of the users have one or more feature values missing). We use the first half of the campaign as training, and the second half for testing. We fit the conversion probabilities $(\theta_{10}, \theta_{11})$ in the training set with probit regressions as done by steps 3–4 of Algorithm 1.

We test the following targeting policies with training data: random, $F(X_i) = 1$, conversion optimization, $\theta_{11}|X_i$, and maximization/minimization of ATE, $\{ATE(X_i), -ATE(X_i)\}$. We also test a variant of the ATE maximization, where the

Table 1: Simulator Validation. Targeting functions are trained and tested with the same data. ATE intervals are shown for 0.10 significance level.

$F_{targ}(X_i)$	**ATE** (1e-6)	*lift* (%)	$F_{targ}(X_i)$	**ATE** (1e-6)	*lift* (%)
1(Random)	3.76±9.83	**7.37**	$\theta_{11}\|X_i$	2.92±10.0	5.46
ATE(X_i)	5.63±9.62	11.77	-ATE(X_i)	-1.74±10.3	-2.94
ATE$^+(X_i)$	8.74±9.53	**19.26**	-ATE$^-(X_i)$	-6.68±10.9	**-9.78**

Table 2: Targeting Policy Testing Results. ATE intervals are shown for 0.10 significance level.

	All Users		No Missing Features	
$F_{targ}(X_i)$	**ATE**(1e-5)	*lift*(%)	**ATE**(1e-5)	*lift*(%)
1 (Random)	1.35±1.74	11.01	2.21±4.26	14.06
$\theta_{11}/(1-\theta_{11})\|X_i$	1.38±1.77	10.91	1.98±3.85	12.25
ATE(X_i)	1.45±1.73	12.00	2.45±4.39	16.25
ATE$^+(X_i)$	1.69±1.76	**13.72**	2.92±3.55	**19.92**
$lift^+(X_i)$	**1.78±1.76**	**14.47**	3.00±3.42	**20.87**

segments with negative ATE are set to the minimum positive ATE ($ATE^+(X_i)$), and likewise for the minimization of ATE ($-ATE^-(X_i)$). Table 1 shows the results. As expected, maximizing ATE shows the best performance, and minimizing ATE the worst ($lift = 19.29\%$ for $ATE^+(X_i)$, and $lift = -9.78\%$ for $-ATE^-(X_i)$). Both estimations are far from the random targeting ($lift = 7.37\%$) validating the simulator. Table 2 shows the testing results. We find that the performance of the user conversion optimization ($\theta_{11}/(1-\theta_{11})$) is similar to that of a random targeting (10.91% versus 11.01%). The best performance is provided by optimizing the lift and setting the negative segments to the minimum positive lift ($lift^+(X_i)$ with 14.47%), which is the only significant effect at 0.10 statistical level (1.78±1.76e-5). We show the effect results estimated for users with no missing features, which depict the same directional results with larger confidence intervals.

4. DISCUSSION AND FUTURE WORK

We have proposed a user targeting simulator that uses data from standard ad effectiveness causal estimation. We have found evidence that the standard practice of optimizing the conversion probability does not optimize the causal effect of the ad. We have shown that the user targeting makes a difference in the ad evaluation even when a placebo ad is displayed. This finding contradicts the standard evaluation practice of measuring the effect with a non-optimized campaign, which is assumed to hold for future optimized exposures. Future directions include the evaluation of behavioral targeting, and the evaluation of ATE-optimized targeting polices from non-experimental data (study group).

5. ACKNOWLEDGMENTS

This work is partially funded by CONACYT UC-MEXUS grant 194880, CITRIS and AOL Faculty Award

6. REFERENCES

[1] J. Barajas, J. Kwon, R. Akella, A. Flores, M. Holtan, and V. Andrei. Marketing campaign evaluation in targeted display advertising. In *ADKDD '12*, pages 1–7. ACM, 2012.

[2] R. A. Lewis, J. M. Rao, and D. H. Reiley. Here, there, and everywhere: correlated online behaviors can lead to overestimates of the effects of advertising. In *Proceedings of WWW Conference 2011*, pages 157–166, 2011.

[3] S. Pandey, M. Aly, A. Bagherjeiran, A. Hatch, P. Ciccolo, A. Ratnaparkhi, and M. Zinkevich. Learning to target: What works for behavioral targeting. In *CIKM '11*, pages 1805–1814, 2011.