

Dynamic Evaluation of Online Display Advertising with Randomized Experiments: An Aggregated Approach

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

Jaimie Kwon
AOL Research
Palo Alto CA, USA
jaimie.kwon@teamaol.com

Aaron Flores
AOL Research
Palo Alto CA, USA
aaron.flores@teamaol.com

Victor Andrei
AOL Research
Palo Alto CA, USA
victor.andrei@teamaol.com

ABSTRACT

We perform a randomized experiment to estimate the effects of a display advertising campaign on online user conversions. We present a time series approach using Dynamic Linear Models to decompose the daily aggregated conversions into seasonal and trend components. We attribute the difference between control and study trends to the campaign. We test the method using two real campaigns run for 28 and 21 days respectively from the *Advertising.com* ad network.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*; G.3 [Mathematics of Computing]: Probability and Statistics—*Time Series Analysis*

Keywords

A/B Testing; Causal Attribution; DLM; Marketing

1. INTRODUCTION

Online Marketing Campaign evaluation has received a great amount of attention by the research community and industry recently. The allocation of a given budget to online display advertising as a marketing channel has motivated the development of statistical methods to measure its effectiveness accurately. The use of randomized experiments, also known as A/B testing in industry, has demonstrated to be effective to evaluate marketing campaigns without over-estimating their effects [4, 2]. These methods require a time window where users are tracked and the measures of interest are collected. As a result, the estimation is aggregated for that time window. This aggregation is a limitation as often sales are affected by seasonal effects. For instance, detecting which days of the week a given campaign is more effective provides more insight to understand and improve the campaign.

We propose a time series approach to estimate the effects of marketing campaigns on the daily number of sales or con-

versions. We consider the randomized design proposed by Barajas *et al.* in targeted display advertising [2]. In the current paper, users are randomized into control and study groups before any decision has been made in the targeting process. We aggregate the daily number of conversions over all the users and consider these sales time series for the control and the study groups. We decompose these series jointly into weekly and trend components using Dynamic Linear Models (DLM) [5]. Based on this framework, we infer the daily mean causal effect as the sales trend differences between both series. In previous work, we developed a method to estimate these effects without randomized experiments, and using the sales trend before the campaign as baseline [1]. In this approach, we incorporate a more accurate baseline, which allows us to draw causal conclusions from the randomized experiment.

2. METHODOLOGY

We define y_t^{ct} and y_t^{st} as the total number of online conversions observed for users in the control and study groups respectively. We assume a latent space model, using a DLM, where a seasonal and trend sales components for both treatment groups are modeled in the state evolution. We define:

$$\begin{aligned} y_t^{ct} &= F^{(0)}\theta_t^{(0)} + F^{(tr)}\theta_t^{ct(tr)} \\ y_t^{st} &= F^{(0)}\theta_t^{(0)} + F^{(tr)}\theta_t^{ct(tr)} + F^{(tr)}\theta_t^{st(tr)} \end{aligned} \quad (1)$$

$\theta_t^{(0)}$ represents the state of a shared (background) base model, which we assume to be a weekly seasonal model. $\theta_t^{ct(tr)}$ is the trend model for the control group, and $\theta_t^{st(tr)}$ is the difference in sales trends attributed to the campaign. $F^{(0)}$ and $F^{(tr)}$ represent observational matrices to model the trend and the base components respectively. We consider the case of unbalanced probabilities of user assignment to the study group $z = 1$, and to the control group $z = 0$, which are fixed from the randomized experiment. We write this model as a 2-D DLM as follows:

$$\begin{aligned} Y_t &= F'\theta_t + \nu_t & \nu_t &\sim N(0, V) \\ \theta_t &= G\theta_{t-1} + w_t & w_t &\sim N(0, W) \end{aligned} \quad (2)$$

where:

$$\begin{aligned} Y_t &= [y_t^{ct}, y_t^{st}]' & \theta_t &= [\theta_t^{ct(tr)}, \theta_t^{st(tr)}, \theta_t^{(0)}]' \\ F' &= \begin{bmatrix} p(z=0) & 0 \\ 0 & p(z=1) \end{bmatrix} \times \begin{bmatrix} F^{(tr)} & 0 & F^{(0)} \\ F^{(tr)} & F^{(tr)} & F^{(0)} \end{bmatrix} \end{aligned} \quad (3)$$

*Main contact.

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

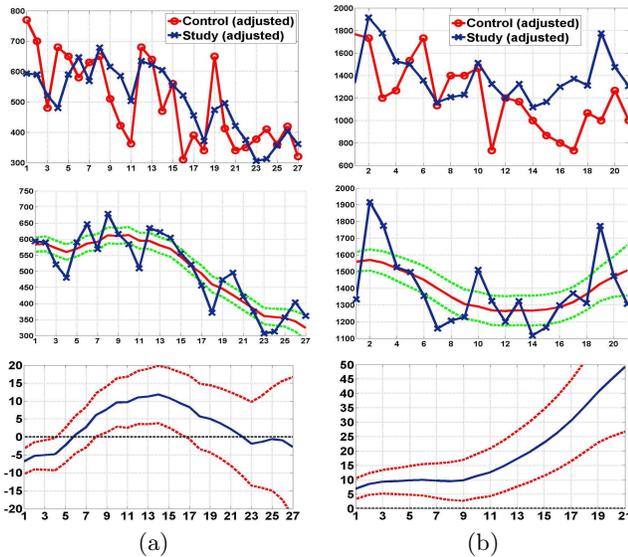


Figure 1: Dynamic Attribution for: (a) campaign 1, (b) and campaign 2. From top to bottom, observed conversions adjusted based on $p(z)$, series trend fitted for the study group, and dynamic causal lift CL_t .

We set $F^{(tr)}$ and $F^{(0)}$ to model a random walk trend and a weekly seasonal components. Similarly, G is constructed as the superposition of the basic components assumed. Thus, these matrices are fixed based on these simpler models¹. We find the Maximum Likelihood (ML) estimates of the variances V, W using an Expectation Maximization (EM) approach [3]. Given V and W , we estimate the distribution of the latent states $P(\theta_t | Y_{1:T})$ for $t = 1, \dots, T$ using the Kalman filtering and backward smoothing equations (E-step). We then optimize the augmented likelihood after replacing the expected values for each state (M-step)². These steps are performed iteratively until convergence.

Given the ML estimates $\{\hat{V}, \hat{W}\}$, we smooth the time series to find the expected causal trend difference attributed to the campaign. We find the causal lift (CL_t) as the percentage change in sales trends, due to the campaign, with respect to the the control trend: $CL_t = 100 \times F^{(tr)} \theta_t^{st(tr)} / F^{(tr)} \theta_t^{ct(tr)}$. We use the Delta method to approximate the distribution of the ratio of two Normal random variables.

3. RESULTS

Fig 1 shows the results for two real campaigns. As illustrated, the attribution is not evident from the observed data. This is a consequence of the seasonal component that affects both series, and typical noisy conversion data. We observe from the causal lift evolution that there are positive and negative effects for campaign 1 at different times. Even when the observed data suggests this positive effect, comparing point by point is highly problematic and it does not provide any statistical support. This behavior shows a campaign with immediate effects where at the beginning of the campaign users wait to buy, probably to survey the competition. Then, the campaign effects peak to gradually fade to the prior-campaign sales level. For campaign 2, positive

¹See [5] pages 89-95 for the random walk trend, and 102-106 for the Fourier seasonal models to fix these components.

²For details of the optimization see [3].

Table 1: Mean attribution lift (%) estimated from the trend differences and the raw data.

Method	Campaign 1			Campaign 2		
	Low	Med	High	Low	Med	High
MCL - Trend	1.31	3.11	4.91	17.03	19.47	21.90
MCL - Raw	-5.03	1.31	7.65	8.29	14.50	20.71

effects are clear from the observed data towards the end of the series. This is why the causal lift shows an increasing tendency. This analysis illustrates that campaign 2 provides delayed effects after the campaign is finished, as opposed to campaign 1. Table 1 shows the average campaign effects estimated from the series trends, and from the raw data. The mean causal lift (MCL - Trend) is obtained as the average CL_t for the campaign duration for both treatment groups. We compare this estimation with the raw estimation (MCL - Raw), obtained from the sample mean of the observed data points without using the time sequence. As depicted, this raw measure is noisier and does not provide any insight about the time when the campaign is more effective.

4. DISCUSSION AND FUTURE WORK

We have presented a time series based approach to attribute trend differences to marketing campaigns. We attribute these differences using causal estimates based on a randomized experiment. We constrain the evolution to be smooth to avoid sudden changes in the attribution. This method provides disaggregated estimates that allows us to obtain marketing insights about the time that the campaign is more effective. The approach we have presented is an aggregated analysis over users. This considers the number of users as unobservable random variables. As on-going work, we will incorporate the series of the total number of users exposed to the campaign and those who are not. We will model these user visitations and exposures as time series in a joint distribution. Conditioning the analysis on the number of users is a common practice. However, this practice is non-trivial when the users convert at a different time than the one they are exposed to the campaign. We will consider these cases with post-treatment variables and survival analysis techniques.

5. ACKNOWLEDGMENTS

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

6. REFERENCES

- [1] J. Barajas, R. Akella, M. Holtan, J. Kwon, and B. Null. Measuring the effectiveness of display advertising: a time series approach. In *WWW (Companion Volume)*, pages 7–8, 2011.
- [2] J. Barajas, J. Kwon, R. Akella, A. Flores, M. Holtan, and V. Andrei. Marketing campaign evaluation in targeted display advertising. In *ADKDD '12: Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*, pages 1–7. ACM, 2012.
- [3] Z. Ghahramani and G. Hinton. Parameter estimation for linear dynamical systems. Technical report, 1996.
- [4] 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 the WWW2011*, pages 157–166, 2011.
- [5] G. Petris, S. Petrone, and P. Campagnoli. *Dynamic Linear Models with R*. use R! Springer-Verlag, 2009.