Toward Improving Digital Attribution Model Accuracy

Stephanie Sapp, Jon Vaver

Google Inc.

Abstract

The accuracy of an attribution model is limited by the assumptions of the model, and the quality and completeness of the data available to the model. Common digital attribution models on the market make a critical, yet hidden, assumption that ads only affect users by directly changing their propensity to convert. These models assume that ad exposure does not change user behavior in other ways, such as driving additional website visits, generating branded searches, or creating awareness and interest in the advertiser. In a previous paper [1], we described a Digital Advertising System Simulation (DASS) for modeling advertising and its impact on user behavior. In this paper, we use this simulation to demonstrate that current models fail to accurately capture the true number of incremental conversions generated by ads that impact user behavior, and introduce an Upstream Data-Driven Attribution (UDDA) model to address this shortcoming. We also demonstrate that development beyond UDDA is still required to address a lack of data completeness, and situations that include highly targeting advertising.

1 Introduction

Digital attribution algorithms use observational user-level path data to assign credit for conversions to the marketing events to which a user was exposed prior to converting. Credits for conversions are typically assigned at the user-level, with a credit value assigned to events within each converting user's path. Overall values for each marketing event type are computed by aggregating across user-level credits. Advertisers use the values assigned to each marketing type to assess the effectiveness of their advertising and make decisions regarding their media spend.

A user's path contains all of the user's observed interactions with a given advertiser, ordered by the time at which each interaction occurred. For example, the user's second observed interaction with the advertiser appears at index two within the user's path. Figure 1 shows an example user path. The user's first observed touch-point with the advertiser is a direct navigation to the advertiser's website (index 1), followed by a display ad impression (index 2). The user later clicks on the advertiser's search ad (index 3), and converts after this paid click.

Position-based attribution models (PBA), also referred to as position rule-based models, are the simplest class of attribution algorithms. This type of model analyzes converting paths only, and assigns credit for conversions deterministically according to the position of marketing events preceding each conversion. The most basic PBA algorithms are "single source" models, which assign all credit for each conversion to only one preceding event. The most common single source PBA attribution model is "last interaction" or "last click", which assigns full credit to the last marketing event that was observed prior to each conversion (the search ad click, in the example user path shown in Figure 1). "First interaction" assigns full credit to the first mar-

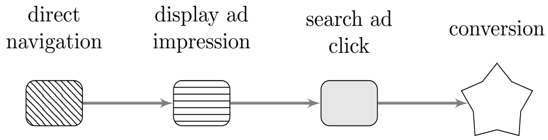

Figure 1: Example user path, consisting of observed interactions with an advertiser. This path includes a conversion, and both paid (display ad impression, search ad click) and organic (direct navigation) touch-points.

keting event observed (the direct navigation, in the example user path). “Fractional” PBA models split credit for conversion among multiple preceding marketing events. “Linear”, for example, divides credit equally among all marketing events observed prior to each conversion ($\frac{1}{3}$ credit each to the direct navigation, display ad impression, and search ad click, in the example user path).

Previous studies have shown that PBA models fail to capture the true causal value of advertising. For example, [2] demonstrated that PBA model estimates differ significantly from randomized controlled experimental results, and [1] used simulation to show that these models do not accurately estimate the true number of incremental conversions generated by advertising across several simple scenarios (see Figures 4–7 of [1]). These results are not surprising, since PBA models are widely recognized to make unrealistic assumptions.

Data-driven attribution (DDA) models, also referred to as algorithmic or probabilistic attribution models, assign credit for conversions in a less prescriptive way, and are advertised as better able to capture the value of advertising. This type of model analyzes both converting and non-converting user paths, and adapts the amount of credit assigned to events preceding a conversion based on the conversion rates observed across these two groups of users. Both “matched-pairs” and “discrete choice” models are types of DDA algorithms.

Matched-pairs DDA algorithms are based on the Shapley value [3], a method from game theory to divide credit among n players in a cooperative game setting [4]. These algorithms assign credit to a specific marketing event within a converting user path by comparing the conversion rate of users with the same observed path, to users whose paths do not contain the specific marketing event but are otherwise the same. More specifically, these models compare groups of users with the same observed events before and after the index of exposure or non-exposure. Figure 2 provides an illustration. The difference between the average conversion rates of these two groups determines the preliminary credit that the marketing event receives. The final assignment of credit for converting paths is determined by repeating the same process described above for each event.

Attribution models typically require the sum of credits within a user path to equal the total number of conversions within that path. Therefore, the event-level credits are normalized to sum to the total number of conversions within each path. By attributing to both paid and unpaid event types, the hope is that credits to unpaid events provide a reasonable estimate for the number of conversions that would have occurred in the absence of advertising.

DDA algorithms based on discrete choice use covariates to estimate consumer choices among a set of discrete alternatives [5], and are typically fitted using logistic or probit regression [6]. Since DDA providers consider their algorithms to be proprietary, detailed descriptions about how these models are applied to attribution are not readily available. Without this transparency, it is impossible to evaluate their efficacy. However, based on typical discrete choice model usage in other domains described in the literature ([7], [8]), it is likely that the discrete user choice of whether (or how often) to convert is modeled as a function of the number of occurrences of each observable marketing event type within a user path, and the impact of each marketing event on conversion is then estimated as

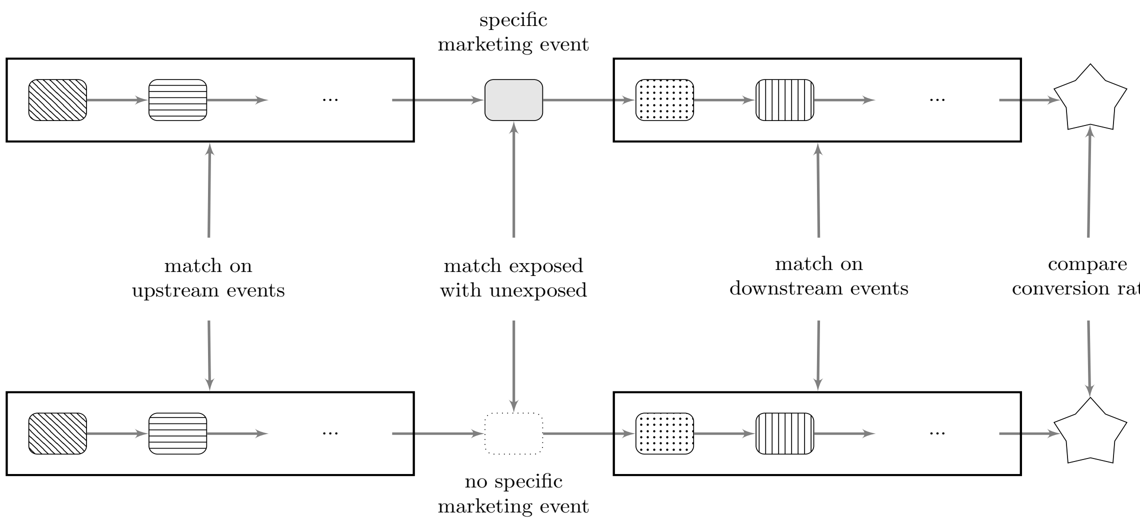

Figure 2: Illustration of typical matched-pairs DDA algorithm. Users with the same observed sequence of events both prior to and after the index of ad exposure are aggregated. A similar unexposed set of users is also aggregated. Credit assigned to the ad exposure using the difference in conversion rates between the exposed and unexposed user sets. These credits are aggregated across all unique exposed user paths.

a function of the coefficients from the regression model.

Additional features are frequently added to the basic DDA models described above. Typical features described in marketing materials of attribution model providers ([9], [10]) include those designed to handle the effects of advertising recency and frequency, ad serving bias, data sparsity, demographics, psychographics, offline factors, and baseline conversion rate impact. While these added features are beyond the scope of this paper, none of them address the key shortcoming of DDA models described here.

In the following sections, we use display ad impressions to illustrate attribution model performance. Note, however, that all descriptions and explanations that use display ads as an example apply to any marketing event type.

2 DDA Shortcoming

2.1 Description

Common DDA models currently offered by attribution providers make the same underlying assumption that the *only* effect of ads is to directly modify a user's probability of conversion. More specifically, these models assume that ads *do not* change a user's downstream *browsing* behavior. For example, the models assume that ads do not increase a user's likelihood of visiting the advertiser's website, or performing a related branded or generic search. Instead, if the user happens to navigate to the advertiser's website, their probability of conversion may be higher or lower as a result of ad exposure.

For matched-pairs DDA models, this underlying assumption is a result of matching on events both upstream and downstream of ad exposure (a display impression, for example) in user paths. This matching implicitly assumes that ad exposure does *not* affect the exposed user's downstream events. Specifically, matching on downstream events assumes that all users exposed to the ad would have had the same sequence of downstream events if they had not been exposed. Otherwise, the exposed and unexposed users do not form a valid comparison group.

With discrete choice DDA models, this assumption is a consequence of ad effects estimated from (functions of) regression coefficients. Effect estimates from regression coefficients are conditional on the remaining covariates in the model being held fixed. Including post-exposure variables in the model (downstream events) invalidates the interpretation of the regression coefficient as causal, since intervention on the exposure variable can also cause changes in the post-exposure variables [11].

2.2 Simulations with DDA

We demonstrate how DDA models fail to capture the value of ads if they change user behavior by evaluating the performance of an example matched-pairs DDA algorithm with a set of simulations generated using the Digital Advertising System Simulation (DASS) [1]. DASS simulates online user browsing behavior, ad serving, and the impact of ad exposure on user behavior. DASS simulates user browsing behavior through an extended, non-stationary Markov model consisting of three components. First, a user activity path model characterizes user browsing behavior in the absence of advertising. Second, an ad serving model describes the process by which users are exposed to advertising events. Third, an ad impact model specifies how exposure to ads impacts downstream user behavior.

DASS has a wide range of capabilities; the most relevant to this paper is its ability to vary the behavioral impact of advertising on users. For example, ads may be configured in the simulator to impact the downstream browsing behavior of exposed users (e.g., ad exposure can increase the likelihood that a user will do a branded search), or may alternatively be configured to only impact the user's likelihood of conversion conditional on site visit, without changing the probability of site visit.

We provide two simulation scenarios to illustrate DDA model performance, which highlight this common ad impact assumption. Both scenarios include a single display ad type, with a small click-through rate. A complete description of the simulation parameters used is provided in Appendix A.

The true number of incremental conversions generated from display ads is calculated by running a "virtual experiment" for each simulation, in which we generate an additional simulated data set with display ads turned off (all other simulation parameters remain the same), and count the number of conversions lost by turning off the display ads (see Section 3 of [1] for more details). The DDA results are generated by applying the DDA algorithm to user paths of observable events types (display ad impressions, conversions, direct navigations, and paid and organic clicks to the advertiser's website) from each simulation. We then compare the conversions attributed to display ads by the DDA algorithm to the true number of incremental conversions determined by running the virtual experiment.

In the first scenario, display ads impact user behavior by increasing the user's likelihood of performing a related branded or generic search, or visiting the advertiser's website. The ad impact parameterization is provided in Appendix A.1. Figure 3(a) shows the true number of incremental conversions from display ads, compared to the number of conversions attributed to display ads by a DDA algorithm, across different levels of display ad effectiveness. If display ads are completely ineffective, the DDA algorithm succeeds in assigning minimal credit to these ads. However, when display ads change user browsing behavior, which also leads to additional conversions, DDA fails to capture that value.

In the second scenario, display ads impact user behavior by directly increasing the user's proba-

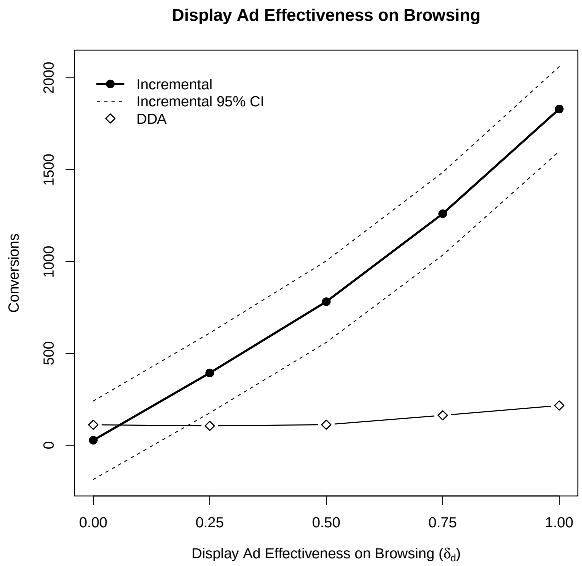

(a) Simulations with varying levels of display ad effectiveness in changing user browsing behavior, by increasing the likelihood of performing a related search, or visiting the advertiser's website. DDA algorithms fail to capture the value of ads that change user behavior in this way.

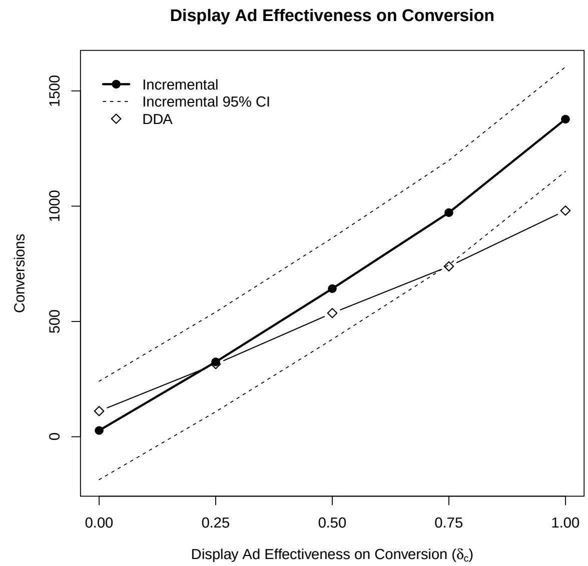

(b) Simulations with varying level of display ad effectiveness in directly changing user conversion rate. The number of credits attributed by DDA to display ads is close to the true number of incremental conversions from display ads.

Figure 3: Scenarios differentiating DDA model performance depending on how ads impact user behavior.

bility of conversion, but only if the user happens to visit the advertiser's website. The ads do not change the user's downstream browsing behavior in any other way. The ad impact parameterization is provided in Appendix A.2. While this mechanism of ad impact may be less applicable for digital advertising, it illustrates the assumption required for DDA models to perform well. The results are shown in Figure 3(b). DDA accurately captures the value of advertising if ads *only* directly impact a user's likelihood of conversion, and do not change browsing behavior in any other way.

3 Upstream DDA

3.1 Algorithm Description

As demonstrated in the previous section, DDA algorithms fail to capture the full impact of advertising if it affects user behavior in ways beyond a direct change in conversion rate. This deficiency is caused by matching user activities downstream from the point of ad exposure, which may have been affected by the ad exposure. To address this problem, this section describes the Upstream DDA (UDDA) algorithm, developed based on the Rubin causal model framework [12]. UDDA estimates ad impact by comparing the conversion rates of user paths with the same sequence of events *upstream* from the index of exposure to the ad, to the conversion rate of unexposed paths that have the same upstream sequence of events, but no ad exposure at the exposure index. Note that this modified matching procedure is similar to the way in which an experiment would operate, since it is only possible and legitimate to control for pre-exposure variables.

Figure 4 shows how UDDA compares similar user paths in its calculations to estimate causal ad impact. Users with an identically ordered set of events prior to the index of ad exposure are aggregated, with no matching on events that occur after the index of ad exposure. The average conversion rate of these paths is compared to the average conversion rate of a second set of paths that are unexposed to the ad, but have the same set of upstream events prior to the index of exposure. Credit is assigned to the ad exposure using the difference in conversion rates between the exposed and unexposed user paths. These credits are aggregated across all unique exposed user upstream paths.

More formally, let user i have an ordered path of observed events denoted by the vector $X^i = (X_1^i, \ldots, X_{L(i)}^i)$, with length $L(i)$. Let A denote an ad impression exposure of the type currently being analyzed, and let C denote a conversion event. For the index m, let $U_m^i = (X_1^i, \ldots, X_{m-1}^i)$ denote the upstream path that includes the first $m-1$ events in a user path, i.e. the sequence of events that occurred prior to index m in the path. The UDDA algorithm then operates as follows:

1. Classify all user paths as either containing or not containing an ad impression exposure A:
   (a) For all exposed paths, i.e. $\exists A \in X^i$, set exposure indicator $T_i = 1$.
   (b) For all other paths, which are unexposed, set $T_i = 0$.
2. For each user path i with $T_i = 1$:
   (a) Let m_i denote the index of the first occurrence of A in the path. That is, $X_{m_i} = A$, and no previous event in the path equals A.
   (b) Record the upstream path of user i as the sequence of events prior to this exposure index: $U_{m_i}^i$.
   (c) Calculate the number of downstream conversions that occur after index m_i as: $C_i = \sum_{j=m_i+1}^{L(i)} I(X_j^i = C)$.
3. For each unique upstream path u_j from Step 2b:
   (a) Let $n_j(T=1)$ denote the number of users in the exposed group with upstream path u_j. Calculate the average conversion rate among these users as: $\bar{c}_j(T=1) = \frac{1}{n_j(T=1)} \sum_{i=1}^{n_j(T=1)} C_i$.
   (b) Find all unexposed users i with upstream path u_j from index m_j, i.e. $U_{m_j}^i = u_j$. Calculate the number of downstream conversions C_i after index m_j for each of these users analogous to Step 2c. Calculate the average conversion rate $\bar{c}_j(T=0)$ among these users analogous to Step 3a.
   (c) Estimate the incremental conversion rate among these users as the difference in conversion rate among exposed versus unexposed users: $\hat{r}_j = \bar{c}_j(T=1) - \bar{c}_j(T=0)$.
   (d) Estimate the number of incremental conversions among these users as: $\hat{r}_j \cdot n_j(T=1)$.
4. Aggregate the estimated incremental number of conversions among all exposed user paths by aggregating over all unique upstream paths: $\hat{\theta} = \sum_{u_j} \hat{r}_j \cdot n_j(T=1)$.

3.2 Simulations with UDDA

We now show results applying the UDDA algorithm to the same sets of simulation scenarios from Section 2. The scenario, in which display

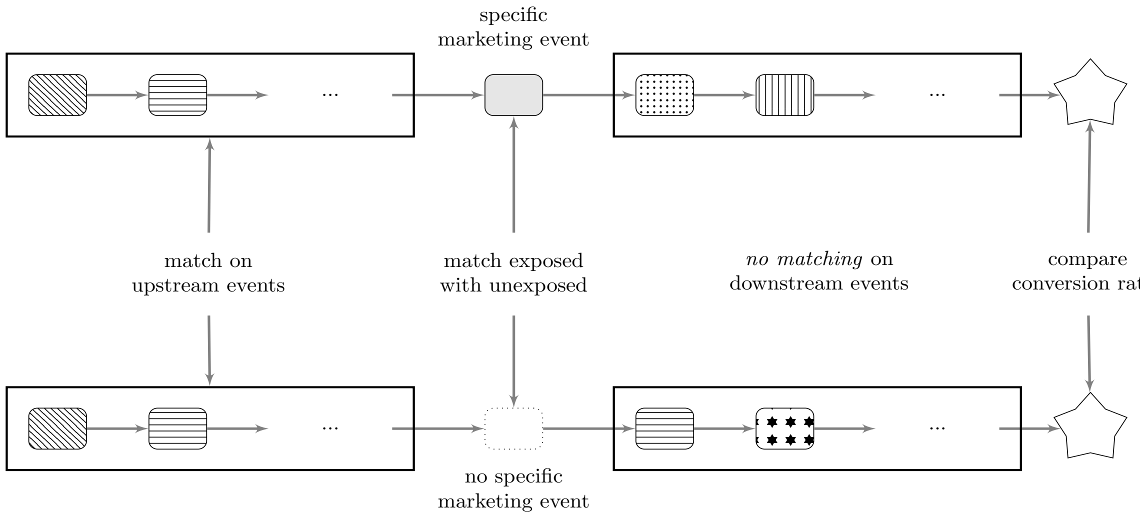

Figure 4: Illustration of UDDA algorithm. Users with the same observed sequence of events prior to the index of ad exposure are aggregated. A similar unexposed set of users is also aggregated. Credit is assigned to the ad exposure using the difference in conversion rates between the exposed and unexposed user sets. These credits are aggregated across all unique exposed user paths. We emphasize that the UDDA algorithm does not perform any matching on events that occur after the index of ad exposure.

ads impact user behavior by increasing the user's likelihood of performing a related branded or generic search, or directly navigating to the advertisers website, is shown in Figure 5(a). The UDDA algorithm better captures the shape of the impact of ads that modify user behavior in this way, and is able to detect an increasing degree of impact from more effective ads. However, note that UDDA misestimates the true incremental conversions by a fixed amount across all levels of effectiveness.

The second scenario, in which display ads impact user behavior by directly increasing the user's probability of conversion if the user happens to visit the advertiser's website, are shown in Figure 5(b). In this case, the UDDA algorithm also follows the shape of the true incremental impact from these ads. As in the previous scenario, UDDA misestimates the true incremental conversions by a fixed amount across all levels of effectiveness. Further, the amount of misestimation is the *same* as in the first scenario. This fixed amount of misestimation is caused by the systematic censoring of users who have no observed interactions with the advertiser. We discuss this issue further in Section 4.1.

4 Ongoing Challenges

This section identifies three remaining challenges for attribution algorithms: systematically censored users, user browsing dissimilarity, and ad targeting. We illustrate how these issues cause problems for the UDDA algorithm, but note that these problems apply to common DDA models, as well.

4.1 Systematically Censored Users

In Section 3.2, we highlighted the observation that the UDDA algorithm misestimates true incremental conversions by a fixed amount for both simulation scenarios and across all levels of ad effectiveness. This misestimation is caused by the systematic censoring of users who have no

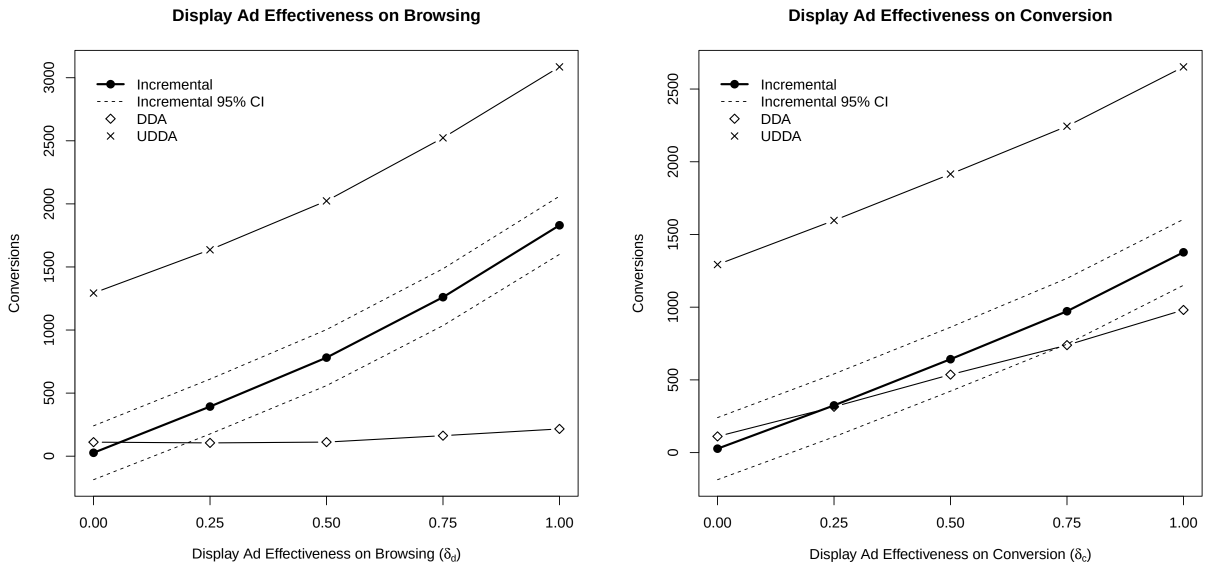

(a) Simulations that vary the level of display ad effectiveness in changing user browsing behavior. UDDA better captures the impact of advertising that changes user downstream browsing behavior, but misestimates by the same fixed amount across all effectiveness levels.

(b) Simulations that vary the level of display ad effectiveness in directly changing user conversion rate. UDDA follows the same shape as the ground truth, but misestimates by the same fixed amount across all effectiveness levels.

Figure 5: Comparison of UDDA and DDA performance with the same scenarios shown in Figure 3.

observed interactions with the advertiser. Since conversion events are always preceded by an observed website visit event, the censoring is not at random, but is instead systematically related to the conversion outcome. This type of missingness is referred to in the literature as missing not at random (MNAR) or nonignorable missingness [13], and introduces estimation bias [14].

Without accounting for these censored users, the UDDA algorithm is unable to accurately estimate the mean conversion rate of all unexposed users. Specifically, recall that the UDDA algorithm estimates ad impact by comparing the conversion rates among all user paths with the same sequence of events upstream from the index of ad exposure, versus paths that have the same upstream sequence of events, but no exposure at the exposure index. When the index of exposure is one (that is, ad exposure is the first observed event in the path), there is no upstream sequence of events on which to match, so the corresponding unexposed user set should consist of all unexposed users. However, we do not have visibility to the users who have no observed interaction with the advertiser, so we are missing the information needed to estimate the conversion rate of all unexposed users (also called the "baseline conversion rate")[1]. Note that this issue only impacts our ability to estimate our ad impact estimates for exposed paths with an exposure index of one. That is, when the exposure index is two or higher, there is necessarily one or more upstream events on which to match, and so we always observe the required set of unexposed users.

To illustrate how lack of visibility to these unobserved users impacts the performance of UDDA, we apply UDDA to the same sets of simulation scenarios presented previously in Section 2 and

[1]In Figure 5(a) and Figure 5(b), we handled this missingness by assigning all credit for conversions in paths beginning with a display ad to the display ads, which is why UDDA overestimates the number of incremental conversions.

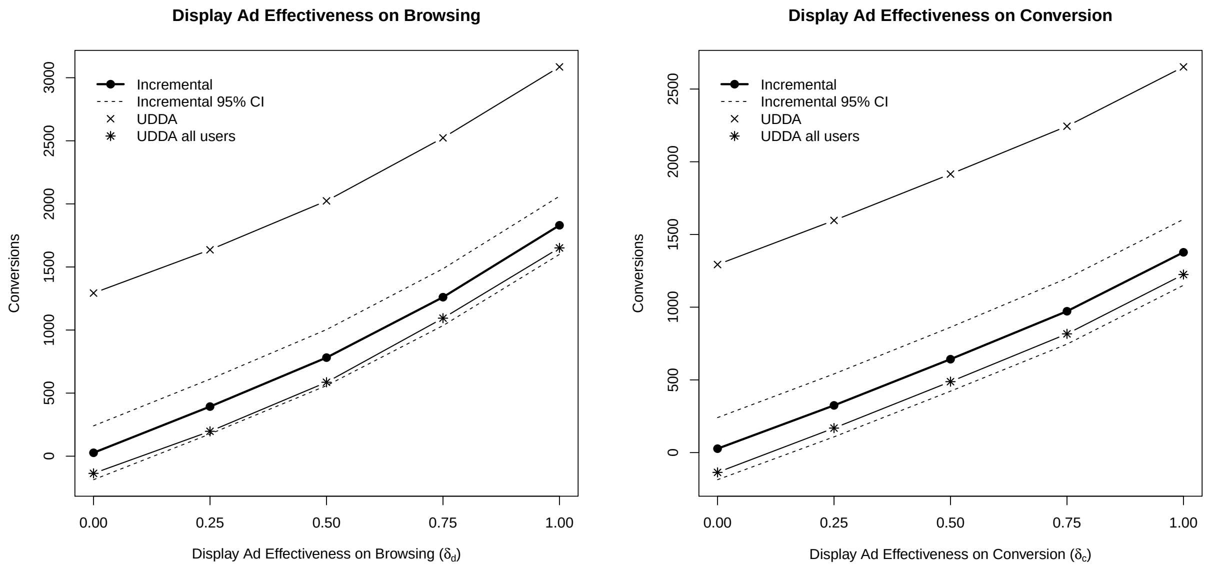

(a) Simulations that vary the level of display ad effectiveness in changing user browsing behavior. UDDA *with visibility to censored users* recovers true incremental conversions.

(b) Simulations that vary the level of display ad effectiveness in directly changing user conversion rate. UDDA *with visibility to censored users* recovers true incremental conversions.

Figure 6: Comparison of UDDA with and without visibility to censored users on the same scenarios shown in Figure 5.

Section 3, except now *with visibility to censored users*[2]. The first simulation scenario, in which display ads impact user browsing behavior, are shown in Figure 6(a), and the second simulation scenario, in which display ads directly impact user conversion probability, are shown in Figure 6(b). In both sets of scenarios, UDDA *with visibility to censored users* does well in estimating the true number of incremental conversions.

Note that visibility to censored users is not possible in practice. However, these results demonstrate the need to account for these users in some way in order to accurately estimate the baseline conversion rate among all users unexposed to the ad type of interest.

[2]DASS simulates activity streams for all users. These streams are then filtered into paths of observable events within the example data scope (Table 3). Visibility to users who are censored due to the data scope is possible in the simulation, since we know the number of users lost during the filtering step.

4.2 User Browsing Dissimilarity

Even in the idealized situation in which we have visibility to censored users, two additional challenges remain which limit our ability to accurately estimate the true incremental conversions generated from advertising. The first challenge is caused by the dissimilarity between the browsing behavior of exposed versus unexposed users. In particular, all users exposed to advertising visited at least one website on which that ad type is eligible to be served. For example, users exposed to a search ad are known to have visited a search engine website. We have no similar knowledge among the unexposed user set. Specifically, among the unexposed users, we cannot identify which ones visited a website on which the ad type was eligible to be served. In general, not all unexposed users will have had the opportunity to be served the ad type because they never engaged in the relevant activity (performing a search, in the case of search ads).

We illustrate how user browsing dissimilarity creates estimation problems by modifying the first simulation scenario. The original version of this simulation specified the initial activity state of all users to be a third-party site visit. Since display ads are also served on this activity state, we ensured that all users had the opportunity to be served this ad type. In the modified version of this simulation, we split the initial activity state among users between search and third-party site visits, with equal probability of starting on either state. Note that search is a more engaged activity than third-party site visit, so the probability of visiting the advertiser's website is higher when starting with a search activity. The parameter modification is provided in Appendix A.3.

Since users who start out with a search activity may never visit a third-party site, the modified simulation includes users who never had the opportunity to be served a display ad. Figure 7 shows the results of this browsing dissimilarity simulation scenario. Even with visibility to censored users, UDDA is unable to accurately estimate the true number of incremental conversions. In this scenario, UDDA underestimates the impact of display ads, since the unexposed user set includes users who were more engaged, and hence more likely to convert, but never had the opportunity to be served a display ad. The conclusion is that even *idealistic* matching on upstream paths is not enough to ensure that two groups of users are comparable to one another.

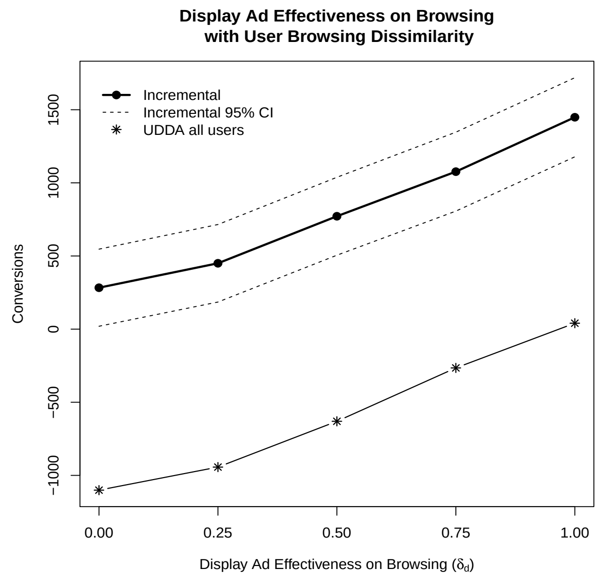

Figure 7: Simulations that vary the level of display ad effectiveness in changing user browsing behavior, modified to include user browsing dissimilarity. Some unexposed users never had opportunity to be served a display ad. UDDA *with visibility to censored users* is unable to recover true incremental conversions.

4.3 Ad Targeting

The second additional estimation issue is caused by targeted advertising. Specifically, ads targeted towards users who behave differently than untargeted users. This type of targeting introduces challenges because attribution models estimate ad impact by comparing conversion rates between exposed versus unexposed users. The rationale is similar to the browsing dissimilarity issue described in Section 4.2. All users exposed to advertising met the ad's targeting criteria. However, we have no similar knowledge among the unexposed user set. That is, we cannot identify which unexposed users met the ad's targeting criteria. In general, not all unexposed users meet the ad's targeting criteria.

Another modification to the first simulation scenario illustrates the problem caused by ad targeting. Note that we resume use of the original initial activity state of all users. That is, we remove browsing dissimilarity by having all users begin their browsing with a third-party site visit, which ensures that all users have the opportunity to be served a display ad. The modification introduced here is the specification of a heterogenous set of users. Specifically, we specify two types of users. The first user type has a higher conversion rate than the second user type, and ads are targeted primarily to the first user type. A full specification of this modification is described in Appendix A.4. Figure 8 shows the results of this

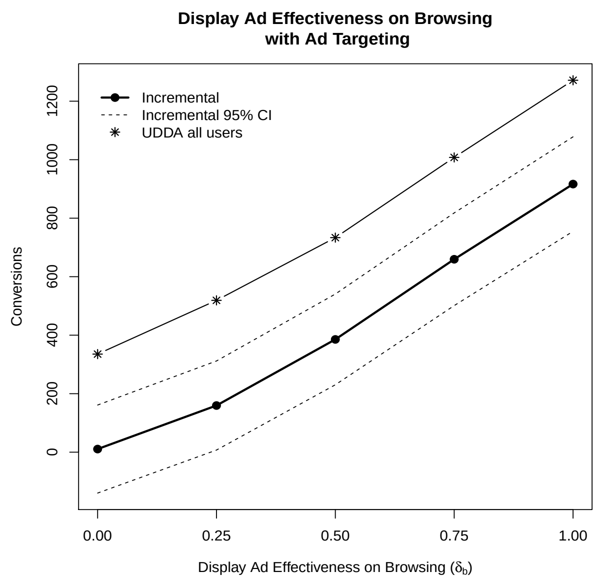

Figure 8: Simulations that vary the level of display ad effectiveness in changing user browsing behavior, modified to include ad targeting. Some unexposed users do not meet the ad's targeting criteria. UDDA *with visibility to censored users* is unable to recover true incremental conversions.

simulation scenario with ad targeting. As in the browsing dissimilarity case, even with visibility to censored users, UDDA is unable to accurately estimate the true number of incremental conversions. In this scenario, UDDA overestimates the impact of display ads, since the unexposed user set includes users who were less likely to convert, but were not targeted by the ad.

5 Concluding Remarks

In this paper, we use simulation to demonstrate that DDA models rely on an implicit assumption that ads do not change the browsing behavior of exposed users, and instead only directly affect users' conversion probability. We introduce the UDDA algorithm, which does not control for post-exposure user activities, and showed it is better able to capture the impact of advertising. Finally, we identified three additional issues: systematic user censoring, user activity dissimilarity, and ad targeting; that limit the performance of attribution models, including UDDA.

We are actively researching ways to address these issues. One approach is the use of additional modeling. For example, modeling the conversion rate among exposed users, and using this model to predict the conversion rate among unexposed users.

Another strategy is to incorporate new data sources. Some ad systems record data which allow us to identify instances in which a user was both active and targeted by an ad, but unexposed. For example, certain auctions log bid results, consisting of the losing participants of each cookie auction. By making this type of data available to an attribution model, a more comparable set of unexposed users can be identified.

In forthcoming papers, we will describe how both additional modeling and the incorporation of new data sources improves attribution model performance.

Acknowledgments

The authors wish to thank the many colleagues from Google whose contributions enhanced this paper, with special thanks to Tony Fagan for his many helpful suggestions.

Appendix A Simulator Parameterization

The simulations in this paper used a simplified version of the parameters as those described in [1]. In particular, we used a simplified set of activity states and only one ad type. For completeness, this appendix fully specifies all simulation parameters used to generate the results in this paper.

Activity State	Description
s	search
vp	visit to a website that the advertiser owns via a click on a paid ad
vup	visit to a website that the advertiser owns via any non-paid click activity
tpw	third party website visit (website not belonging to the advertiser)
c	conversion
eos	end of session

Table 1: Description of activity states used in simulations.

The total number of users is set to $K = 10,000,000$. The set of activity states is

$$a_1, \ldots, a_n = \{\text{s}, \text{vp}, \text{vup}, \text{tpw}, \text{c}, \text{eos}\}$$

where the definition of each state is provided in Table 1.

We set the initial activity state distribution so that all users begin with a third party website visit, the activity tpw. That is, we set $\pi_{a_i} = 1$ for $a_i = \text{twp}$, and $\pi_{a_i} = 0$ for $a_i \neq \text{tpw}$.

The transition matrix M, consisting of probabilities p_{a_i,a_j} is specified as:

$$M = \begin{bmatrix} & \text{s} & \text{vp} & \text{vup} & \text{tpw} & \text{c} & \text{eos} \\ \text{s} & .08 & 0 & .03 & .33 & 0 & .56 \\ \text{vp} & .08 & 0 & .03 & .33 & .03 & .53 \\ \text{vup} & .08 & 0 & .03 & .33 & .03 & .53 \\ \text{tpw} & .08 & 0 & .01 & .33 & 0 & .58 \\ \text{c} & .08 & 0 & .03 & .33 & 0 & .56 \\ \text{eos} & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

We use one display ad type in the simulations:

$$b_1, \ldots, b_m = \{b_1\} = \{\text{dsp}\} = \text{display ads}$$

The associated ad serving parameters for the display ad type b_1 is shown in Table 2.

Each user's impressibility to the display ads is determined by random draw from the following truncated normal distribution:

$$q_{b_1}^k \sim Tr[0,1]N(\mu = 0.8, \sigma = 0.1)$$

Ad Type b_j	Serving States s_{b_j}	Impress Thresh d_{b_j}	Share of Voice v_{b_j}	Freq Cap f_{b_j}
dsp	tpw	0.8	0.4	100

Table 2: Ad serving parameters used in simulations.

The example data scope of observable event types is shown in Table 3.

Observable Event	Description
organic search click	unpaid visit to the advertiser's website is immediately preceded by a search activity state
display ad impression	display ad impression is served
other non-ad visit	unpaid visit to the advertiser's website is immediately preceded by a non-search activity
conversion	conversion activity state is reached

Table 3: Summary of observable events within our example data scope.

To parameterize display ad response, we use two functions which apply different time horizon effects. The two timeframe effect functions are: temporary impression $r_{\text{ti}}^{\text{dsp}}$ and persistent impression $r_{\text{pi}}^{\text{dsp}}$. Specification of these functions differs depending on how the ads impact user behavior. Appendix A.1 describes the function specification when the ads change user downstream browsing behavior. Appendix A.2 describes the function specification when the ads directly impact user conversion probability.

The function $\hat{f}(n_k^{b_1})$ appears as part of the temporary and persistent impression effect functions. This function serves to modify the impact of an ad based on the number of times $n_k^{b_1}$ a user was exposed to the ad. Specifically, columns that are scaled up increase in effect over the first several impressions (burn-in), and then decline in effect after additional impressions (fatigue). The following ad burn-in / fatigue correction function is used to modify the magnitude of changes to the user's transition matrix due to the $n_k^{b_1}$-th display ad exposure:

$$\hat{f}(n_k^{b_1}) = \begin{cases} \frac{n_k^{b_1}}{n_0} & n_k^{b_1} \leq n_0 \\ 2 - \frac{n_k^{b_1}}{n_0} & n_0 < n_k^{b_1} < 2n_0 \\ 0 & n_k^{b_1} \geq 2n_0 \end{cases}$$

In this hat function, n_0 specifies the ad exposure that results in the maximum ad impact. For the example simulations presented in this paper, we set $n_0 = 2$.

A.1 Display Ads Impact Browsing

When display ads impact the downstream browsing behavior of users, the temporary impression effect function is given by:

$$r_{\mathrm{ti}}(\delta_b) = h_{\mathrm{nrm}} \circ h_{\mathrm{sc}}^{\mathrm{s,vup},1.2\delta_b \hat{f}(n_k)} \circ h_{\mathrm{spk}}^{\mathrm{vp},0.001}(M)$$

The function sets transition probabilities of the vp column to 0.001 (to model the possibility of a paid ad click), scales the transition probabilities in the columns s and vup by a factor of $1.2\delta_b \hat{f}(n_k^{b_1})$ (increasing the likelihood of a visit to the advertiser's website, or performing a related search), and then re-normalizes the transition matrix.

The persistent impression effect function for display ads is given by:

$$r_{\mathrm{pi}}(\delta_b) = h_{\mathrm{nrm}}^{\mathrm{c}} \circ h_{\mathrm{sc}}^{\mathrm{s,vup},1.2\delta_b \hat{f}(n_k)}(M)$$

This function scales the transition probabilities in the columns s and vup by a factor of $1.2\delta_b \hat{f}(n_k^{b_1})$, and then re-normalizes the matrix. The re-normalization is performed using all columns except the conversion column c, in order to keep the conversion rate given site visit constant.

A.2 Display Ads Impact Conversion

When display ads directly impact user conversion probability, the temporary impression effect function is given by:

$$r_{\mathrm{ti}}(\delta_c) = h_{\mathrm{nrm}} \circ h_{\mathrm{spk}}^{\mathrm{vp},0.001}(M)$$

The function sets transition probabilities of the vp column to 0.001, and then re-normalizes the transition matrix.

The persistent impression effect function for display ads is given by:

$$r_{\mathrm{pi}}(\delta_c) = h_{\mathrm{sc}}^{\mathrm{c},1.2\delta_c \hat{f}(n_k)}(M)$$

This function scales the transition probabilities in the columns c by a factor of $1.2\delta_c \hat{f}(n_k^{b_1})$ (directly increasing the likelihood of conversion).

A.3 Browsing Dissimilarity Modification

User browsing dissimilarity is added by a modification of the initial activity state distribution so that users have an equal probability of beginning with either a third party website visit (tpw) or a search (s). That is, we set $\pi_{a_i} = 0.5$ for $a_i \in \{\text{twp, s}\}$, and $\pi_{a_i} = 0$ for $a_i \notin \{\text{tpw, s}\}$.

A.4 Ad Targeting Modification

We introduce ad targeting through specification of a heterogenous set of users. Specifically, we model two different types of users, with an equal number of users of each type. The first user type is specified with the same transition matrix and impressibility to display ads as described previously in Appendix A, and repeated here for convenience:

$$M_A = \begin{array}{cc} & \begin{array}{cccccc} \text{s} & \text{vp} & \text{vup} & \text{tpw} & \text{c} & \text{eos} \end{array} \\ \begin{array}{c} \text{s} \\ \text{vp} \\ \text{vup} \\ \text{tpw} \\ \text{c} \\ \text{eos} \end{array} & \begin{bmatrix} .08 & 0 & .03 & .33 & 0 & .56 \\ .08 & 0 & .03 & .33 & .03 & .53 \\ .08 & 0 & .03 & .33 & .03 & .53 \\ .08 & 0 & .01 & .33 & 0 & .58 \\ .08 & 0 & .03 & .33 & 0 & .56 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{array}$$

$$q_{b_1}^{k_A} \sim Tr[0,1]N(\mu = 0.8, \sigma = 0.1)$$

The second user type has a lower inherent conversion rate, and impressibility drawn from a distribution with a lower mean:

$$M_B = \begin{array}{cc} & \begin{array}{cccccc} \text{s} & \text{vp} & \text{vup} & \text{tpw} & \text{c} & \text{eos} \end{array} \\ \begin{array}{c} \text{s} \\ \text{vp} \\ \text{vup} \\ \text{tpw} \\ \text{c} \\ \text{eos} \end{array} & \begin{bmatrix} .08 & 0 & .03 & .33 & 0 & .56 \\ .08 & 0 & .03 & .33 & .01 & .55 \\ .08 & 0 & .03 & .33 & .01 & .55 \\ .08 & 0 & .01 & .33 & 0 & .58 \\ .08 & 0 & .03 & .33 & 0 & .56 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{array}$$

$$q_{b_1}^{k_B} \sim Tr[0,1]N(\mu = 0.6, \sigma = 0.1)$$

Ad serving parameters remain the same. Specifically, the impressibility threshold for display ads remains 0.8. As a result, ads are targeted primarily towards the first user type.

References

[1] Stephanie Sapp, Jon Vaver, Minghui Shi, and Neil Bathia. "DASS: Digital Advertising System Simulation" *http://research.google.com/pubs/pub45331.html*, 2016

[2] David Reiley, Randall Lewis, and Taylor Schreiner. "Ad Attributes and Attribution: Large-Scale Field Experiments Measure Online Customer Acquisition" *http://ssrn.com/abstract=2049457*, 2012

[3] Lloyd S. Shapley. "A Value for n-Person Games" *Annals of Mathematical Studies 28:307–317*, 1953

[4] Alvin E. Roth. "The Shapley Value: Essays in Honor of Lloyd S. Shapley" *Cambridge University Press*, 1988

[5] Moshe E. Ben-Akiva and Steven R. Lerman. "Discrete Choice Analysis: Theory and Application to Travel Demand" *Vol. 9. MIT Press*, 1985

[6] Kenneth Train. "Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand" *Vol 10. MIT Press*, 1986

[7] Daniel McFadden and Kenneth Train. "Mixed MNL Models for Discrete Response" *Journal of Applied Econometrics 15(5):447–470*, 2000

[8] Kenneth Train. "Recreation Demand Models with Taste Differences Over People" *Land Economics 74(2):230–239*, 1998

[9] Robert Stratton. "Which Analytical Method Best Meets the Attribution Challenge?" *MarketShare: A Neustar Solution*, 2016

[10] Xiaolong Zhang. "Cross-Channel Response Attribution: Credit where Credit is Due" *Experian Marketing Services*, 2013

[11] Andrew Gelman and Jennifer Hill. "Data Analysis Using Regression and Multilevel/Hierarchical Models" *Cambridge University Press*, 2006

[12] Donald B. Rubin. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies" *Journal of Educational Psychology 66.5: 688*, 1974

[13] Roderick J.A. Little and Donald B. Rubin. "Statistical Analysis with Missing Data" *John Wiley & Sons*, 2014

[14] Yiran Dong and Chao-Ying Joanne Peng. "Principled Missing Data Methods for Researchers" *SpringerPlus 2.1*, 2013