Off-Policy Evaluation of Probabilistic Identity Data in Lookalike Modeling^{*†}

Randell Cotta, Mingyang Hu, Dan Jiang and Peizhou Liao

Drawbridge Inc.

San Mateo, California

{randy,mingyang,dan.jiang,peizhou.liao}@drawbridge.com

ABSTRACT

We evaluate the impact of probabilistically-constructed digital identity data collected between Sep. 2017 and Dec. 2017, approximately, in the context of Lookalike-targeted campaigns. The backbone of this study is a large set of probabilistically-constructed “identities”, represented as small bags of cookies and mobile ad identifiers with associated metadata, that are likely all owned by the same underlying user. The identity data allows to generate “identity-based”, rather than “identifier-based”, user models, giving a fuller picture of the interests of the users underlying the identifiers.

We employ off-policy evaluation techniques to evaluate the potential of identity-powered lookalike models without incurring the risk of allowing untested models to direct large amounts of ad spend or the large cost of performing A/B tests. We add to historical work on off-policy evaluation by noting a significant type of “finite-sample bias” that occurs for studies combining modestly-sized datasets and evaluation metrics based on ratios involving rare events (e.g., conversions). We illustrate this bias using a simulation study that later informs the handling of inverse propensity weights in our analyses on real data.

We demonstrate significant lift in identity-powered lookalikes versus an identity-ignorant baseline: on average ~70% lift in conversion rate, CVR, with a concordant drop in cost-per-acquisition, CPA. This rises to factors of $\sim (4-32)\times$ for identifiers having little data themselves, but that can be inferred to belong to users with substantial data to aggregate across identifiers. This implies that identity-powered user modeling is especially important in the context of identifiers having very short lifespans (i.e., frequently churned cookies). Our work motivates and informs the use of probabilistically-constructed digital identities in the marketing context. It also deepens the canon of examples in which off-policy learning has been employed to evaluate the complex systems of the internet economy.

^{*} All authors contributed equally to this work.

[†] Although the data used in this work is of a proprietary nature and will not be released publicly, Drawbridge Inc. is open to potential work with outside academic researchers. Those interested can reach out to randy@drawbridge.com

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

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

© 2019 Copyright held by the owner/author(s).

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

<https://doi.org/10.1145/3289600.3291033>

KEYWORDS

ACM proceedings, Marketing and Advertising, Identity Management, Entity Resolution, User Modeling and Recommendation Systems, Off-Policy Evaluation.

ACM Reference Format:

Randell Cotta, Mingyang Hu, Dan Jiang and Peizhou Liao. 2019. Off-Policy Evaluation of Probabilistic Identity Data in Lookalike Modeling. In *The Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19)*, February 11–15, 2019, Melbourne, VIC, Australia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3289600.3291033>

1 INTRODUCTION

The portion of the world economy devoted to marketing has grown much faster than the global economy itself. This is largely due to the rise of digital connectivity and the subsequent application of technologies that optimize the matching of brand messages with consumers. While powerful in principle, a proliferation of devices and maturation of web services has led to a complex and fractured view of consumers’ interests, making it difficult to optimally match messages with consumers in practice.

Many of the largest companies in the world have become large precisely by solving this problem. They pair compelling free services with terms-of-use that allow the collection and monetization (via advertising) of vast amounts of rich personal data on large pools of users. Such use of personal data obviously has privacy implications and securing and guaranteeing the proper use of such data has proven challenging in practice [3, 24, 29]. Given the tension between i) the need for companies to effectively reach customers and ii) the need to allow consumers options in terms of how their data is collected and used, a significant opportunity exists in solving this problem outside of this small set of massive platforms.

The synthesis of *Probabilistic Digital Identity Data* is one alternative solution to this problem. Here the idea is that sources of relatively “thin” data coming mainly from the digital advertising ecosystem¹ can be used to build inferred “users” as collections of identifiers². This organization can then be used to estimate the interests/behaviors of the “users” who can, in turn, be messaged via their associated identifiers. As the name suggests, probabilistic digital identity data is created by leveraging statistical machine learning and data mining techniques, a route which is much more technically difficult compared to the deterministic aggregation of user-volunteered data around long lived identifiers (e.g., customerID). The data used to build the probabilistic identities in this work was

¹ This data can be seen as a massive stream of relatively uninformative event tuples (discussed later).

² By which we mean reset-able and often short-lived marketing identifiers like cookies and mobile ad-IDs. We refer to these simply as “identifiers” in what follows.

passively collected and did not contain any information that would directly identify an individual such as name, phone number, or postal address. This kind of data is much less rich than the kinds of information typically volunteered to OSNs, thus, on the one hand, potentially lowering the risks associated with data breach/misuse, while on the other hand necessitating significant modeling to predict even basic demographic estimates (age, gender, etc.).

While the usefulness of this data is conceptually intuitive to marketers, the impact of this data is often not straightforward to measure in practice. There are several different problem scopes for which one can consider evaluating the impact. There is also the issue of simply obtaining accurate estimates of effect size in the context of such complex systems. Likely owing to these factors, the authors know of no other work that has previously attempted to carry out these measurements. We attempt to do this here.

In terms of scope, we focus our evaluation of the impact of identity data on the important case of lookalike-targeted campaigns. Lookalike-targeting refers to the continual estimation of user interests and subsequent messaging of users that closely resemble those who have previously resonated with a given message. This is an interesting place to focus not only because of the strategy's general importance in marketing, but also because its analysis poses an intermediate level of difficulty³. To be concrete then, we are interested in answering the following question: *If we allowed identity-powered user models to drive our lookalike targeting, instead of identity-ignorant variants, what would be the change in performance of the system?* By identity-ignorant lookalike-targeting we refer to systems based on user models that estimate the interests of the underlying user using only the data that is associated to a single identifier. Identity-powered systems, in contrast, model users' interests by aggregating data across all of the identifiers that are inferred to belong to the same underlying user.

To answer this question we lean heavily on off-policy evaluation [41]. The reasons for this are two-fold. First, we would like to avoid a) the risk of putting untested models into production systems that control large amounts of ad spend and b) the costs of performing many A/B tests to tune/validate new models. To the latter point, although A/B tests remain the gold standard for estimating the causal impact of interventions in principle, many have noted the often-impractical nature of carrying out these kinds of tests in practice (e.g., [21, 26, 32]). Second, given that we will evaluate these models "off-line", using only data that was previously collected by an "on-line" model, it is important to correct a naive analysis to account for the fact that the naive results are generally not the same as would occur if the new model were actually allowed to make its own decisions and collect its own data.

In the remainder of this paper, we will give a description of the raw data and the probabilistic processing that produces the identity data whose application we would like to assess the impact of. We will then discuss and analyze identity-powered lookalike targeting, using off-policy learning techniques to arrive at a reliable comparison of identity-powered and identity-ignorant systems. Finally, we discuss potential extensions of this work and conclude.

³The importance of identity in re-targeting campaigns is obvious, but the analysis is not as interesting (amounting to a counting exercise), while an analysis for very broadly-targeted Brand campaigns is probably too difficult to do convincingly here due to the lack of signals to quantify the effects.

2 PROBABILISTIC DIGITAL IDENTITY DATA

The basic aim in generating probabilistic identity data is to organize massive sets of identifiers into small clusters such that the identifiers within a cluster are likely associated to the same human user. **Note in particular that the system cannot identify the devices or any other facts associated to one particular individual, but rather, it simply attempts to determine that all identifiers in the cluster are likely owned by one person, and not two or more people.** From a technical perspective, the process involves aspects of information retrieval (what pairs of identifiers are even relevant to consider?), supervised learning (how likely is it that these two identifiers belong to the same person?), unsupervised learning (detecting users within a complex network of scored pairs of identifiers) and user-modeling (given a user cluster, what can we say about the user's interests?).

We first discuss the raw data upon which probabilistic identity data will be built, and then continue to discuss the processing logic that actually builds it. Further detail can be found in [17].

2.1 Raw Data

The raw data used to build identities comes from the platforms that power the digital advertising ecosystem. While this kind of data is massive in scale⁴, it comes as a stream of relatively uninformative, passively collected events. The signal-to-noise in such data falls well short of that typically volunteered by users to platform services like social networks, requiring significant processing to uncover even a basic knowledge of the underlying users and their interactions.

The information inside of these events can be classified into two broad categories: i) observational data, typically including an identifier, timestamp and network address and ii) semantic data, describing the content context of the event (typically an associated content category⁵, demographic estimates for the user, etc.). The identifiers associated to each event tuple are the cookies and/or mobile advertising identifiers used in the advertising industry to allow brands to reach consumers with their messaging. They are user-resettable unique identifiers (generated as large strings by cryptographic hashing algorithms). Active identifiers may thus become inactive at any time, new identifiers may spawn at any time, and many identifiers have a lifespan of only a single observation.

Given this kind of input data it is not even obvious which identifiers are associated to the same device, much less to the same user or to users living in the same household. We now discuss a probabilistic system that discovers this kind of identity structure from this unstructured raw data.

2.2 Creating Identity Data

The system that estimates the underlying identity structure from batches of raw data events consists of a chain of modules that are repeatedly applied in each run to build up hierarchical levels of organization (devices as collections of identifiers, users as collections of devices, households as collections of users). The basic components of the chain are: i) pair discovery, ii) pair scoring and iii) community detection.

⁴There are more than 100B of such events occurring daily across the globe.

⁵Like: entertainment->movies

2.2.1 Pair Discovery. The pair discovery step in the system is an information-retrieval step used simply to control the scale of what is processed by more expensive steps further downstream. Pair discovery performs a TF-IDF-like [34] scoring of identifier co-occurrences on bipartite identifier and “proxy” networks, where the various proxy objects correspond to spatio-temporal localizations, for example, an (IP address, date) tuple. A conservative cut on this score allows the most relevant $\sim 100\text{B}$ pairs of identifiers to be considered further.

2.2.2 Pair Scoring. In the pair scoring step supervised learning algorithms are trained to score the pairs produced by pair discovery according to obfuscated-deterministic labels. These labels are constructed using identifier-handle linkages in which the handles are obfuscated and typically refer to something like a hashed email. We label pairs in which the same handle is linked to both identifiers as +1’s, and those for which no handle is linked to both identifiers as -1 ’s⁶. Given this, one can imagine that +1’s are pairs for which someone can verify that the same email has been used on both identifiers and -1 ’s are pairs for which we have failed to verify the same. We do expect that our -1 ’s are erroneous at some rate simply due to the imperfect knowledge of our label data providers and attempt to source identifier-handle linkages from multiple providers to reduce label errors that occur due to limited visibility.

Pair scoring, as a supervised learning task in exactly this context, has been the subject of prior work termed as “cross-device” linkage scoring. To date there have been multiple machine learning competitions [1, 2] regarding techniques for estimating which ids likely belong to the same underlying user [7, 12, 23, 25, 30, 33, 40, 44, 46, 48], among other works [39, 45, 47]. The features used in most pair scoring models, include those in the system described here, include several classes: i) those regarding the behaviors of each of the two objects in the pair in isolation ii) those having to do with the patterns of interaction between these two objects and iii) a featurization of the graph contexts in which this pair of objects is embedded⁷. While features in ii) typically do the real heavy lifting, features in iii) are particularly important for accurately scoring pairs that do not already have a long history of interaction. While there are a handful of models involved in producing the features themselves, once produced we utilize an ensemble of boosted [14] and bagged [6] trees to produce the final scores.

2.2.3 Community Detection. If one views the set of scored pairs as a graph, this graph will be a web-scale complex network of weighted undirected edges. The graph has an approximately power-law degree distribution, small diameter and degree assortativity as in many other datasets [35]. In particular, simply removing edges with scores below some value will not reduce the graph into one in which the connected components are plausible representations of the true underlying users. For this we must apply a community detection (CD) algorithm [13] to the graph of scored edges.

Given the scale of our graph, $\sim 10\text{B}$ vertices and $\sim 100\text{B}$ edges, the set of CD algorithms and implementations that can be effectively used is highly constrained. Perhaps the most scalable (and therefore

most often used) massive CD algorithm is Label Propagation, LP [42]. The extreme scalability of LP comes from its basis on a simple heuristic rule to create communities, however, and relying on such a rule has several downsides in our context: i) the basic rule leads to low-quality communities as measured both by eye and by truth data; ii) changing the basic rule is possible, but inconvenient as a mapping between tweaks to the rule and the resulting changes in performance is complex (one must reason about rules that may lend toward some particular objective, check if they, in fact, do improve this objective, etc.); iii) found communities are highly unstable to small input graph changes (clusters have little overlap in consecutive runs of the system).

To overcome these issues the system that produces the identity data described here implements a highly-distributed local Fitness optimization procedure for its CD step that we term GFDC (Generalized-Fitness-Driven Clustering). The algorithm is related to many previous works [4, 9, 27, 28, 43], though differing substantially in the definition of “Fitness” that is used. Historically, “Fitness” has tended to refer to the well-known network modularity [36], or to slight variants with free parameters [27, 28] or by addition of other local network structural information [11]. For our purposes, with a network built on noisy and incompletely-observed links, pure modularity optimization does not always produce clusters that plausibly represent underlying users. For GFDC we suppose only that Fitness is a weighted sum of locally-calculable terms describing proposed communities. We use the Fitness function as a way to not only drive toward higher modularity, but also, for instance, to enforce desired behaviors like temporal smoothness [8] or to provide Bayesian encouragement based on domain knowledge. As an example of the latter case, we can add a term that describes high-fitness user clusters as those in which there are about one (rarely zero or two) phones. Such a term would appear alongside a modularity term so that clusters having ten phones that might be proposed by modularity optimization alone would be broken down by the two-term definition of “Fitness” (provided the loss of modularity in doing so is not too great).

The system uses Apache Giraph⁸ to run GFDC at the required scale. The calculation proceeds in rounds, wherein each round a random selection of vertices can decide to join any adjacent community based on a calculation of the changes in Fitness that would result. The decision is greedy (vertices always choose the community that maximizes fitness) and myopic (vertices assume that no other vertices will move, which is not the case). Despite this, the flexibility of our definition of Fitness results in both greatly increased precision at all recall values with respect to an LP benchmark (measured by pairwise labels described in subsection 2.2.2) and clusters that are much more reasonable-looking and stable. Of course, GFDC is also more expensive to run than LP, but given our distributed implementation, not prohibitively so, even at this scale.

2.3 Probabilistic Identity Dataset

Given that probabilistic techniques are being employed on noisy, incomplete, massive scale, dynamic data, we expect that any given set of putative “identities” are afflicted with both Type-I errors (two identifiers placed in the same cluster which are not actually

⁶Given that both identifiers have *some* handle/s attached. We label pairs in which one/both identifiers have no handles as 0’s (this class composing the vast majority of pairs in the system).

⁷For graph embedding methods, see a recent review [16] and the references therein.

⁸<http://giraph.apache.org/>

Table 1: Statistics for some of the Identity datasets used in this work. G_1 (G_8) is the operating point exhibiting the highest tolerance for FPs (FNs) that we consider here. Basic filters have been applied to remove clusters that are not plausible regular consumers (i.e., perhaps bots, etc.). Reported are figures for the number of users internationally and in the U.S. market specifically as well as the median number of identifiers per user and physical devices per user (US).

Dataset	Global (M)	US (M)	ID/Usr (US)	Dev./Usr (US)
G_1	901	264	33	3
G_8	602	176	5	1

associated with the same user) and Type-II errors (two identifiers that are placed in two different clusters, even though they are actually associated to the same user). Any practical implementation of this technology allows tuning between the two ends of this error spectrum (small/tight clusters with very few False Positives (FP) and many False Negatives (FN) vs. large/fuzzy clusters with very few FNs and many FPs) to address use-cases with varying degrees of tolerance to the two types of error. In our system, after users are produced in the CD step, we give each identifier a cluster membership score based on simple measures of affinity calculated on the input network. Cutting on this score “tunes” our user clusters for a given FP/FN tolerance. In order to gauge the effect of this choice on the measurements reported below we present results for a handful of operating points along this spectrum of errors.

The identity datasets that we will use here have been created based on events generated during the period Sep. 2017 - Dec. 2017. After a basic cleaning to remove irregular-looking clusters the whole set is of size ~900M identities, composed of a total of ~23B advertising identifiers that were active at some time during this period. Basic statistics describing our high-FP and high-FN identity dataset endpoints are given in Table 1.

3 IDENTITY-POWERED LOOKALIKES

In this section, we first discuss lookalike systems in more detail. We discuss the metrics that we will use to evaluate performance, data preparation procedures specific to this analysis and features/models employed to model users. We then turn to a detailed description of our off-policy evaluation methodology. We define basic concepts, lay out a simplified model of the larger system that allows to isolate the impact of identity and, finally, discuss a novel kind of bias that arises in the presence of rare-events and finite datasets.

3.1 System, Models, Metrics and Data

The job of a lookalike targeting system is to continually execute: i) collect historical impressions and conversions associated with a given campaign, ii) learn to discriminate between previously converted users and random users, iii) rank non-converted users according to their similarity with the previously converted group and iv) expose highly-ranked users to the campaign message by bidding appropriately to serve them impressions. As this loop is repeatedly iterated the pool of previously converted users grows

Table 2: Basic statistics of the campaign data used here. Quoted for each campaign (Cmp) are the impression scale (Imp), conversion rate for IDs in and out of the production whitelist (CVR_{in} and CVR_{out}, resp.), and descriptions of the campaign market vertical and call to action.

Cmp	Imp(M)	CVR _{in/out} (%)	Vertical	Action
A	1.21	0.049/0.016	Home Security	Engagement
B	7.27	0.014/0.010	Home Goods	Form Fill
C	8.18	0.032/0.009	Arts/Crafts	Purchase
D	0.80	0.053/0.009	Wireless	Form Fill
E	1.46	0.018/0.009	Streaming Video	Form Fill
F	0.45	0.027/0.008	Solar	Land'g Pg.
G	3.71	0.046/0.005	Live Music	Purchase
H	5.43	0.008/0.005	Real Estate	Form Fill

and the system continually refines its ability to message users that particularly resonate with the campaign.

In our system 10M highly-ranked identifiers are whitelisted for rather aggressive bidding, while all identifiers are also potentially impressed in less specifically-targeted channels. This is important in the context of off-policy evaluation: **all** identifiers have some chance to be impressed, so the support in covariate space is not completely collapsed onto that of the whitelisted group. We use gradient-boosted tree ensembles[14, 22] to learn to discriminate previously-converted users from random users. The features we use are primarily semantic features based on the content associated with the user’s events and the basic activity patterns of the identifiers composing the “user”.

Albeit basic in terms of machine learning algorithm and features, the aforementioned system serves as a qualified testing ground for the purpose of evaluating our identity data: The only difference between the identity-powered and identity-ignorant models is whether or not the same set of features is aggregated across all of the identifiers inferred to be associated to the same underlying user. Moreover, due to the generic nature of our system, we expect that the effectiveness of identity data found here would generalize well to other similar but more sophisticated targeting systems.

We focus on CVR as our measure of the performance difference between the two models as we observe similar results (both qualitatively and quantitatively) for CPA as we do for CVR. We filter the data such that only the first conversions and impressions occurring before the first conversion are counted for a given brand/user history, avoiding potential confusion caused by multiple conversions.

We chose a set of eight campaigns to include in the study, which could be useful in understanding the heterogeneity of our results. While the brands involved will remain anonymous, we note that the campaigns were chosen based on several criteria: i) impression and/or conversion scale large enough to reasonably study, though ranging from small to large within that criterion, ii) broad targeting (geographical region, demographics, etc.), and iii) collectively representing a variety of products and calls to action. Basic statistics describing the campaign data are provided in Table 2.

In all analyses below the identity data used was created based on events that were generated before these campaigns were run. This

Figure 1: Approximate factorization of our system bidding logic as concerns the Lookalike Model. We study a direct change to the variable L , which indirectly changes the variable B . Both R and W remain approximately unchanged.

eliminates the possibility of any leakage of the campaign results via the structure of the identity data itself.

3.2 Off-Policy Evaluation

Naively, to gauge the impact of identity on user-modeling, one simply trains two lookalike models (identity-powered and identity-ignorant versions) in an offline fashion, makes predictions about which users will convert, waits for the production (on-line) system to collect new data and compares their ability to predict which users did eventually convert. This does not answer our real question of interest: if we were to use either of these models *in place of the production system* (part of which includes its own on-line lookalike model that does not use identity data), how would the results differ? New models would evaluate users, adjust bids and serve impressions differently than the online model. To overcome this bias one should re-weight the data collected by the production system to bring estimates into closer accord with what would be expected from alternative systems.

3.2.1 Inverse-Propensity Weighting. A procedure which is commonly used to do this correction is *importance weighting* [19] or *inverse-propensity weighting* [18] (“IPW”, in the case where propensity models are being used to work in high-dimensional covariate spaces). In equations this can be seen as a change of measure on covariate space:

$$Y = \int_{\vec{x}} y P(\vec{x}) \quad \text{and} \quad Y^* = \int_{\vec{x}} y P^*(\vec{x}) = \int_{\vec{x}} y \left(\frac{P^*(\vec{x})}{P(\vec{x})} \right) P(\vec{x}), \quad (1)$$

where Y and Y^* are the outcome quantity of interest (CVR) estimated under the on-line or off-line policy, respectively (Y^* being the main subject of our interest here). $\vec{x}$ is an example point in covariate space. $P(\vec{x})$ and $P^*(\vec{x})$ are the sampling distributions representing systems leveraging the on-line and off-line models, respectively. For our purposes, $P(\vec{x})$ will be defined as the probability that an ad request represented by covariates $\vec{x}$ is impressed under the on-line bidding policy. The estimates for Y and Y^* is then

$$\hat{Y} = \frac{1}{n} \sum_{i \sim P} y_i \quad \text{and} \quad \hat{Y}^* = \frac{1}{n} \sum_{i \sim P} \left(\frac{P^*(\vec{x}_i)}{P(\vec{x}_i)} \right) y_i = \frac{1}{n} \sum_{i \sim P} w_i y_i, \quad (2)$$

where $\sum_{i \sim P}$ means summing over data points i that are sampled from distribution P . Then, we see that estimating the quantity of interest $\hat{Y}^*$ requires up- or down-weighting examples collected under the on-line policy according to weights $w = P^*(\vec{x})/P(\vec{x})$. Examples that would be more commonly seen if the off-line model had been used to collect data rather than the on-line model that was actually used, i.e., when $P^*(\vec{x}) > P(\vec{x})$, are thus up-weighted (and vice-versa).

Given that we are changing only one piece in a larger complex system, one can usefully employ a structural model [37, 49] to factorize the larger $P(\vec{x})$ and $P^*(\vec{x})$ distributions such that many of the factors cancel in their ratio. In our bidder, the influence of the Lookalike model can be isolated using the following Markov factorization (corresponding structural model pictured in Fig. 1):

$$P(\vec{x}) = P(R, L, B, W) \quad (3)$$

$$\approx P(W|B, R) \cdot P(B|L, R) \cdot P(L|R) \cdot P(R) \quad (4)$$

Both of the factors $P(R)$ and $P(W|B, R)$ will cancel out in the ratio, corresponding to the statements that i) The characteristics of requests that pass basic targeting criteria for some campaign (“qualified requests”) is independent of the internal model logic that we will apply in our system and ii) conditioned on the request R and the bid amount B , the external auction is isolated from knowledge of what model we applied to come up with that bid amount. Both of these statements are approximate, but expected to hold quite well. The former is true up to changes in the nature of internal competition between the various campaigns in our own system, and the latter is true up to changes in the auction environment as external competitors adjust to changes in our bidding behavior.

In contrast to $P(R)$ and $P(W|B, R)$, we do not expect the factors $P(B|L, R)$ and $P(L|R)$ to cancel out in the ratio. The factor $P(L|R)$ is precisely what we are changing when we change the Lookalike model, it is the probability of a given user being placed on the Lookalike whitelist with respect to some campaign. Getting on this list depends primarily on the model score (the part we will change by using identity-powered features rather than identity-ignorant features), but also secondarily on seemingly un-involved covariates regarding user activity level and behavior patterns⁹. As there are a handful of such covariates, we will employ simple logistic regression models to estimate $P(L|S, X)$, $P^*(L|S^*, X)$, where $X \subset R$ are activity/behavioral covariates and S and S^* are identity-ignorant and identity-powered Lookalike scores, respectively.

The bid-amount term, $P(B|L, R)$, is also interesting to consider. In our system bid amounts are largely influenced by feedback of outcome values (i.e., bids tomorrow are a function of CVR and CPA today, among other considerations), so, strictly speaking, there is a feedback loop that should be included in our structural model which would invalidate its use in causal studies [38]. Given this, we perform two analyses, where we set the ratio of these factors as:

$$\text{i) } \frac{P^*(B|L, R)}{P(B|L, R)} = 1 \quad \text{and,} \quad \text{ii) } \frac{P^*(B|L, R)}{P(B|L, R)} = \frac{P^*(L|S^*, X)}{P(L|S, X)}. \quad (5)$$

For the first case we set the ratio = 1 (canceling the bid term), which will give estimates for the performance that we should expect in the time before which feedback is incorporated¹⁰, or in the case where feedback is not the primary driver of bid prices. In the second case, we make the following series of assumptions: that the ratio of the bid amounts would scale with the ratio of CVRs for the two models, which itself would scale with the ratio of the two propensity scores. Of course, both of these choices are rough approximations, but useful inasmuch as we expect the first to under-estimate and the second to over-estimate the true steady-state IPWs. We thus

⁹We do not consider all possible users for Lookalike whitelisting, the considered set partly depends on how active the users are, etc.

¹⁰In this case, this is $\sim$ days, as conversion feedback is sparse and latent data.

Figure 2: The sampling distributions P and P^* in the simulation study are shown on the left. The histogram of the resulting IPW is displayed in the center panel, which qualitatively matches the histogram from real data shown on the right. Conversions are plotted as red bars on the x-axis in both histograms, showing the regions in which they are rare.

expect the true performance of a deployed Lookalike model to lie somewhere between the two estimates obtained in what follows.

An important consideration in off-policy evaluation is in how to treat examples falling in regions of support non-overlap, i.e., examples near which the IPW, w , explodes. In such regions a naive weighting by the IPW implies that a few examples dominate the estimation, greatly increasing the variance of our estimates. Typical heuristic solutions trade off most of the variance at the cost of some bias. Here we use weight truncation [31], $w_T(w) = \min(w, w_0)$, where weights are truncated as they cross a threshold w_0 . While many other works [15, 20] have focused on comparing a wide variety of counterfactual estimators (e.g., weight dropping, truncation, model-based and doubly-robust variants) we opted not to do this here, as our main focus is on the evaluation of the impact of identity data, and as we were not able to follow up with the kind of extensive A/B testing that would be required to definitively judge the relative performance of various estimators. We did do our analysis twice, once using the weight truncation mentioned above, and once using weight dropping [5], where weights are set to zero upon crossing a threshold w_0 , finding estimates that were consistent with w_T and of similar variance (these are omitted for clarity in what follows).

3.2.2 Finite-Sample Bias. Besides the bias/variance trade-off that usually determines w_0 , in our analysis we encountered another kind of bias. The issue here is caused by a confluence of factors: the sparse nature of conversions, the modestly-sized (in terms of conversions) campaign datasets used here, and the importance of the upper tail of the IPW distribution. Given the sharply falling tail of the IPW distribution, for some region $w_0 > w_s$ one statistically expects only a small amount of impressions and thus very few conversions, even though the CVR in that regime may well be quite high. This regime of $w_0 > w_s$ is pertinent because it is precisely where we hope the new model to improve upon the old, yet we could be counting thousands of highly-weighted impressions and the estimated CVR collapses to unreasonably low values. We refer to this source of bias as *finite-sample bias*.

A simple simulation study illustrates this effect. It is also a nice sandbox in which practitioners can learn basic rules for setting w_0 in analyses on real data. Here we suppose a population of identifiers indexed simply by $x \in [0, 1]$. Each member of the population has an inherent CVR, $y(x)$, which is linear in x for simplicity. We suppose that the production and new models have true propensity functions

Figure 3: Top, CVR^* as a function of IPW threshold w_0 , is calculated (squares) from two batches of datasets of 50K (green) and 500K samples (blue), and compared to infinite-sample $\text{CVR}^*(w_0)$ (solid black) and ground truth $\text{CVR}^*(\infty)$ (i.e. no truncation, black dotted). Bottom, the total error (solid color) is decomposed into the finite-sample bias (spheres), truncation bias (solid black) and standard deviation (dash-dot). The heuristic choice of w_0 , Eqn. 7, is shown as vertical bars on the x-axis.

of $P(x)$ and $P^*(x)$, respectively, that are Beta distributions (illustrated in Fig. 2), with parameters chosen to match the qualitative behaviors of our actual data, especially in the high-IPW tail¹¹.

Given this setup we can estimate CVR^* both analytically (c.f. Eqn. 1) and numerically (c.f. Eqn. 2), allowing a knowledge of infinite-sample and finite-sample results, respectively, and thus a decomposition of the overall bias, Bias_{all} , into truncation, Bias_T , and finite-sample, Bias_{FS} , components as following:

$$\begin{aligned} \text{Bias}_{\text{all}} &= \widehat{\text{CVR}}^*(w_0) - \text{CVR}^*(\infty) \\ \text{Bias}_T &= \text{CVR}^*(\infty) - \text{CVR}^*(w_0) \\ \text{Bias}_{\text{FS}} &= \text{Bias}_{\text{all}} - \text{Bias}_T \end{aligned} \quad (6)$$

In other words, the overall bias is the difference between the estimated $\widehat{\text{CVR}}^*(w_0)$ and the ground truth $\text{CVR}^*(\infty)$; the infinite-sample truncation bias equals the ground truth minus the true CVR^* with truncation, *analytically* calculated by integrating Eqn. 1 over the range $[0, w_0(x)]$; the finite-sample bias is the leftover bias after accounting for the truncation.

The one term left in Eqn. 6, $\widehat{\text{CVR}}^*(w_0)$, needs to be numerically calculated. Here, 1) one dataset i with 50K or 500K impressed identifiers is sampled from the known Beta distribution $P(x)$. 2) It is then bootstrapped 100 times, allowing us to calculate the median and variance of $\widehat{\text{CVR}}^*$ at different w_0 for *this dataset* i in the same way a normal practitioner would do. Note that, without the knowledge of the underlying probability distribution P , the estimates $\widehat{\text{CVR}}^*(w_0)_i$ when $w_0 \gg 1$ would not be robust due to the relatively large variance in the modest-sized data, and the variances $\text{Var}(w_0)_i$ would be underestimated from bootstrapping[10]. In order to obtain a robust and accurate estimate of $\widehat{\text{CVR}}^*(w_0)$ for the bias breakdown

¹¹The actual distributions were: $P(x) \sim B(2.5, 2.5)$ and $P^*(x) \sim B(0.5, 4.5)$.

Figure 4: Estimates for the lift in CVR for lookalike-targeted campaigns due to the use of identity data. In the left panel, results by-campaign and averaged over campaigns for the G_1 (high-FP) identity dataset. On the right, campaign-averaged results for a series of identity datasets ranging between the G_1 and G_8 endpoints. Naive estimates are presented along with corrected results using overall weights of w_T or w_T^2 , according to the choice of bid-factor (c.f. Eqn. 5).

(Eqn. 6), and to correctly quantify the variance so that the total error $(\text{Bias}_{\text{all}}^2 + \text{Var})^{\frac{1}{2}}$ can be measured, step 1) and 2) are repeated 100 times, and the median of these $\widehat{\text{CVR}}^*(w_0)_i$ and $\text{Var}(w_0)_i$ are used as $\widehat{\text{CVR}}^*(w_0)$ and its variance.

The results from the simulation study are shown in Fig. 3. We observe the expected trade-off between truncation bias at low w_0 and variance at higher w_0 . At even higher w_0 there is the regime with large finite-sample bias (larger for smaller datasets, as expected). Practitioners with real datasets may expect to see an estimator whose variance rises continually with w_0 and whose mean plateaus, but this is not the case: one observes the variance rising until there are no more conversions with higher IPW, around $w_0 \approx 100$, at which point $\widehat{\text{CVR}}^*$ collapses (with compulsory decrease in variance) and eventually falls significantly below the ground truth. Obviously, the optimal choice for w_0 is to minimize the total error shown in the lower panel in Fig. 3. But this is not visible to a practitioner with a single real dataset and no knowledge of the underlying model. We, therefore, look for a simple heuristic that yields approximately optimal results in our toy model.

This choice should obviously depend on both the campaign size and nominal CVR. We find the following heuristic to provide good results in simulation studies:

$$\Pr(k_{>w_0} \leq 5; N_{>w_0}, \bar{Y}) = 1 - \delta, \quad (7)$$

where $k_{>w_0}$ is the number of conversions and $N_{>w_0}$ the number of impressions falling in the IPW region $w > w_0$, $\bar{Y}$ is the nominal CVR, computed as the average in the production system. The number of conversions thus follows a binomial distribution with parameters $N_{>w_0}$ and $\bar{Y}$ and we set $\delta = 0.05$ so that we start truncating weights where we expect ≤ 5 conversions at IPWs this extreme with 95% probability. This choice is illustrated in Fig. 3 with vertical bars on the x-axis, which does a reasonable job at minimizing the total error for datasets of both sizes. We use this heuristic to choose w_0 in what follows.

3.3 Results

The basic results of our lookalike analysis are presented in Fig. 4, where we present campaign-specific and campaign-averaged results

for the high-FP endpoint identity dataset G_1 as well as campaign-averaged results for identity datasets interpolating between our G_1 (high-FP) and G_8 (high-FN) endpoints.

Though the variance of our estimates is not negligible, we observe (Fig. 4, left) that the identity-powered lookalike system appears to show a significant $\sim 70\%$ lift on average¹². There are some campaigns for which the naive and corrected results largely agree. We find that these are generally those in which the identity-powered and identity-ignorant lookalike whitelists have significant overlap. In most cases where naive and corrected results significantly disagree the corrected results are larger, evidently showing that regions in covariate space that are preferred by the identity-powered model, are preferentially those with higher CVR.

Results comparing different identity datasets are also interesting (Fig. 4, right). Here we have some logical expectations for the shape of the lift curve: at very high FNs we should have very small clusters and the lift relative to identity-ignorant approaches should approach 1 (no lift); moving in the opposite direction we expect lift to rise as useful information is included, and then degrade once clusters become so large that the information being included is mostly noise. We indeed observe an increase in relative lift from nearly 1 (no lift) at the G_8 endpoint to a plateau with a higher lift as we approach the G_1 endpoint here, suggesting that our G_1 dataset, although admitting the most FPs of any here, has not reached the point of significant degradation. Unfortunately, our system does not produce identity datasets with higher FP rates than G_1 , which we plan to investigate in the future.

An even more interesting finding is presented in Fig. 5. Here we show that the advantage of the identity-powered system over the identity-ignorant system is indeed much higher where it is logically expected to be: when the observations on each *identifier* are infrequent, but where more information from the other identifiers associated to the same *user* can supplement. To produce Fig. 5 we define a boolean filter that is a good proxy for these statements:

$$f_i(q_{\text{cut}}) = (\text{obs}_i < Q_{\text{obs}}(1 - q_{\text{cut}})) \wedge (\text{cls}_i > Q_{\text{cls}}(q_{\text{cut}})), \quad (8)$$

where obs_i and cls_i are the number of observations on identifier i and the cluster size of the associated cluster, Q_{obs} and Q_{cls} are

¹²This is a number bracketed by the w_T - and w_T^2 -style off-policy corrected estimates (c.f. Eqn. 5), averaged across campaigns. The Naive estimate of the same had lift $\sim 30\%$.

Figure 5: CVR ratio of identity-powered to identity-ignorant lookalike targeting for identifiers with $f_i(q_{\text{cut}})=1$ as a function of q_{cut} (c.f. Eqn. 8). Presented are Naive and w_T or w_T^2 corrected results. Insets illustrate the limiting cases of low and high q_{cut} , where most information is contained on few long-lived identifiers and where information is spread across a variety of short-lived identifiers, resp. Results are averaged over campaigns and identity dataset G_1 is used.

quantile functions for observations and user cluster sizes¹³. $f_i(q_{\text{cut}})$ can thus be used as a boolean filter for including identifier i given a quantile cut q_{cut} . From Fig. 5, we observe that the CVR lift increases strongly with q_{cut} , reaching values in the $\sim(4-32)\times$ range for the extreme case $q_{\text{cut}} = 90\%$ (amounting to $\sim 1\%$ of identifiers in our data). The implication is that identity data is especially important in the limit of frequently-churned identifiers.

4 DISCUSSION

In this work, we investigated the effectiveness of probabilistically-constructed digital identity data in lookalike-targeted ad campaigns. This data takes the form of hierarchical collections of identifiers, represents identities at a variety of resolutions such as physical devices and users, and are mined from a stream of unstructured and thin event tuples using probabilistic machine learning approaches. We employed off-policy evaluation techniques and a careful handling of the bias that occurs in finite-sized datasets, in order to accurately estimate the lift due to identity data in identity-powered systems above the baseline performance of identity-ignorant variants. Our results suggest a significant lift overall and lend guidance in terms of tuning the identity data being used for this purpose. The lift due to using identity data appears particularly large where it logically should: in the limit where observations associated to an underlying user are spread across a wide number of identifiers that must be inferred to belong to that user (via identity data).

There are many directions in which this work could be fruitfully extended. Perhaps most obvious is that it would be good to compare these results to those coming from online A/B tests. Unfortunately,

this was not possible here because the advertising platform collecting the campaign data became unavailable soon after the data was collected. It would also be interesting to bolster the analysis by including many more campaigns and presenting a fuller analysis of the heterogeneity of the results across campaigns. Here we did not observe a consistent pattern with respect to campaign-specific factors (e.g., call-to-action-type, etc.) among the eight campaigns studied.

ACKNOWLEDGMENTS

The authors thank Azad Jacobs, Mike Murphy, Shadi Rasheed, Kamakshi Sivaramakrishnan and Jenny Zhao for helpful discussions.

REFERENCES

- [1] 2015. ICDM 2015: Drawbridge Cross-Device Connections. (Aug. 2015). <http://www.kaggle.com/c/icdm-2015-drawbridge-cross-device-connections>
- [2] 2016. CIKM Cup 2016 Track 1: Cross-Device Entity Linking Challenge. (Dec. 2016). <http://competitions.codalab.org/competitions/11171>
- [3] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. 2016. The economics of privacy. *Journal of Economic Literature* 54, 2 (2016), 442–92.
- [4] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.
- [5] L. Bottou, J. Peters, J. Quiñero-Candela, D. X. Charles, D. M. Chickering, E. Portugaly, D. Ray, P. Simard, and E. Snellson. 2012. Counterfactual Reasoning and Learning Systems. *ArXiv e-prints* (Sept. 2012). arXiv:cs.LG/1209.2355
- [6] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [7] Xuezhao Cao, Weiye Huang, and Yong Yu. 2015. Recovering cross-device connections via mining IP footprints with ensemble learning. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1681–1686.
- [8] Yun Chi, Xiaodan Song, Dengyong Zhou, Koji Hino, and Belle L Tseng. 2007. Evolutionary spectral clustering by incorporating temporal smoothness. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 153–162.
- [9] Aaron Clauset, Mark EJ Newman, and Christopher Moore. 2004. Finding community structure in very large networks. *Physical review E* 70, 6 (2004), 066111.
- [10] A. DasGupta. 2008. *Asymptotic Theory of Statistics and Probability*. Springer New York.
- [11] Fabricio Olivetti de França and Guilherme Palermo Coelho. 2015. A flexible fitness function for community detection in complex networks. In *Complex Networks VI*. Springer, 1–12.
- [12] Roberto Diaz-Morales. 2015. Cross-device tracking: Matching devices and cookies. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1699–1704.
- [13] Santo Fortunato. 2010. Community detection in graphs. *Physics reports* 486, 3–5 (2010), 75–174.
- [14] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [15] Alexandre Gilotte, Clément Calauzènes, Thomas Nedelec, Alexandre Abraham, and Simon Dollé. 2018. Offline A/B testing for Recommender Systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. ACM, 198–206.
- [16] Palash Goyal and Emilio Ferrara. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* 151 (2018), 78–94.
- [17] Devin Guan, Xiang Li, Randy Cotta, Rahul Bafna, and Obuli Venkatesan. 2016. System to group internet devices based upon device usage. (Dec. 6 2016). US Patent 9,514,248.
- [18] Keisuke Hirano, Guido W Imbens, and Geert Ridder. 2003. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71, 4 (2003), 1161–1189.
- [19] Daniel G Horvitz and Donovan J Thompson. 1952. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* 47, 260 (1952), 663–685.
- [20] Nan Jiang and Lihong Li. 2016. Doubly Robust Off-policy Value Evaluation for Reinforcement Learning. In *International Conference on Machine Learning*. 652–661.
- [21] Thorsten Joachims and Adith Swaminathan. 2016. Counterfactual Evaluation and Learning for Search, Recommendation and Ad Placement. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16)*. ACM, New York, NY, USA, 1199–1201.
- [22] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A highly efficient gradient boosting

¹³For reference, $\sim 98\%$ of previously-converted users and $\sim 87\%$ of users considered for Lookalike whitelisting map to non-trivial (multi-vertex) clusters in our base identity dataset G_1 .

- decision tree. In *Advances in Neural Information Processing Systems*. 3149–3157.
- [23] Girma Kejela and Chunming Rong. 2015. Cross-device consumer identification. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1687–1689.
 - [24] Jay P Kesan, Carol M Hayes, and Masooda N Bashir. 2015. A comprehensive empirical study of data privacy, trust, and consumer autonomy. *Ind. LJ* 91 (2015), 267.
 - [25] Michael Sungjun Kim, Jiwei Liu, Xiaozhou Wang, and Wei Yang. 2015. Connecting devices to cookies via filtering, feature engineering, and boosting. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1690–1694.
 - [26] Ron Kohavi, Alex Deng, Brian Frasca, Roger Longbotham, Toby Walker, and Ya Xu. 2012. Trustworthy online controlled experiments: Five puzzling outcomes explained. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 786–794.
 - [27] Andrea Lancichinetti, Santo Fortunato, and János Kertész. 2009. Detecting the overlapping and hierarchical community structure in complex networks. *New Journal of Physics* 11, 3 (2009), 033015.
 - [28] Andrea Lancichinetti, Filippo Radicchi, José J Ramasco, and Santo Fortunato. 2011. Finding statistically significant communities in networks. *PloS one* 6, 4 (2011), e18961.
 - [29] Susan Landau. 2015. Control use of data to protect privacy. *Science* 347, 6221 (2015), 504–506.
 - [30] Mark Landry, Robert Chong, et al. 2015. Multi-layer classification: ICDM 2015 drawbridge cross-device connections competition. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1695–1698.
 - [31] Lihong Li. 2015. Offline evaluation and optimization for interactive systems. (2015).
 - [32] Lihong Li, Jin Young Kim, and Imed Zitouni. 2015. Toward predicting the outcome of an A/B experiment for search relevance. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*. ACM, 37–46.
 - [33] Jianxun Lian and Xing Xie. 2016. Cross-device user matching based on massive browse logs: The runner-up solution for the 2016 CIKM Cup. *arXiv preprint arXiv:1610.03928* (2016).
 - [34] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
 - [35] Mark EJ Newman. 2003. The structure and function of complex networks. *SIAM review* 45, 2 (2003), 167–256.
 - [36] Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the national academy of sciences* 103, 23 (2006), 8577–8582.
 - [37] Judea Pearl. 1998. Graphs, causality, and structural equation models. *Sociological Methods & Research* 27, 2 (1998), 226–284.
 - [38] Judea Pearl. 2009. *Causality*. Cambridge university press.
 - [39] Minh C Phan, Aixin Sun, and Yi Tay. 2017. Cross-device user linking: url, session, visiting time, and device-log embedding. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 933–936.
 - [40] Minh C Phan, Yi Tay, and Tuan-Anh Nguyen Pham. 2016. Cross device matching for online advertising with neural feature ensembles: First place solution at CIKM Cup 2016. *arXiv preprint arXiv:1610.07119* (2016).
 - [41] Doina Precup, Richard S Sutton, and Satinder P Singh. 2000. Eligibility Traces for Off-Policy Policy Evaluation. In *ICML*. Citeseer, 759–766.
 - [42] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. 2007. Near linear time algorithm to detect community structures in large-scale networks. *Physical review E* 76, 3 (2007), 036106.
 - [43] Satu Elisa Schaeffer. 2005. Stochastic local clustering for massive graphs. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 354–360.
 - [44] Lars Ropeid Selsaas, Bikash Agrawal, Chumming Rong, and Thomasz Wiktorski. 2015. AFFM: Auto feature engineering in field-aware factorization machines for predictive analytics. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1705–1709.
 - [45] Ugo Tanielian, Anne-Marie Tusch, and Flavian Vasile. 2018. Siamese Cookie Embedding Networks for Cross-Device User Matching. *arXiv preprint arXiv:1803.10450* (2018).
 - [46] Nam Khanh Tran. 2016. Classification and Learning-to-rank Approaches for Cross-Device Matching at CIKM Cup 2016. *arXiv preprint arXiv:1612.07117* (2016).
 - [47] Elena Volkova. 2017. *Cross-device tracking with machine learning*. Master's thesis.
 - [48] Jeremy Walthers. 2015. Learning to rank for cross-device identification. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*. IEEE, 1710–1712.
 - [49] Sewall Wright. 1921. Correlation and causation. *Journal of agricultural research* 20, 7 (1921), 557–585.