PPiTTA - Preserving Privacy in TV Targeted Advertising

Tzachy Reinman[1,3] and Erez Waisbard[2,3]
[1] *The Hebrew University of Jerusalem, Jerusalem, Israel*
[2] *Bar Ilan University, Ramat Gan, Israel*
[3] *NDS Technologies Ltd., Jerusalem, Israel*

Keywords: Privacy, Targeted Advertising.

Abstract: Targeted advertising involves using a person's personal data to determine the most promising commercials to show that person. While the benefits are clear, the price paid in terms of loss of privacy may be high. In this work we bridge what seems at first to be contradicting requirements – the ability to personalize data and the need to maintain privacy, especially while reporting back the impressions to the advertiser. We provide two schemes that achieve this, each in its own adversarial model. We put an emphasis on modern TV systems and describe the architecture for supporting it.

1 INTRODUCTION

Advertising is "fuelling" the content industry as users are paying for the content by watching the interspersed commercials. In the past, commercials were broadcast to all TV viewers and the cost of an advertisement was proportional to its length and its level of exposure. However, due to the limitations of broadcasting technology, many people were forced to watch commercials in which they had no interest. Nowadays, technology allows us to deliver *different* advertisements to different viewers watching the same show. This gives advertisers a powerful way to address the relevant audience. In addition, measurement systems are able to report back exactly which ads were viewed by each user (denoted *impression*), giving the advertiser a more accurate measure of the relevant exposure of each ad.

Targeted advertising delivers advertisements to an audience that is likely to have an interest in the advertised product. This is done by gathering as much information as possible about each user and then determining the most appropriate target audience for each commercial. Practically, this is actually accomplished in the reverse direction – from a target set of potential commercials, we pick the one which is most appropriate for each user. This selection may be based on demographic data, or information inferred from the viewing habits of users, for example, determining that a user is a sports fan based on the many hours spent tuned into sports channels.

Many users are unhappy about the collection of personal information for advertising ("OUT-LAW News" of September 30, 2009 reported that "US web users reject behavioral advertising "). As a result both regulators and technology providers are acting to ensure that users have the ability to protect their privacy, if they choose to do so. This is already enacted in the EU's Data Protection Directive, which requires users to be notified about any data that is being collected. This directive calls for maximizing the privacy of the user when handling such data and giving the user an option to opt out. This shift toward enhancing users' privacy can also be seen in Google's announcement to include the 'Do Not Track' button in the Chrome browser.

While enhancing privacy is mostly good, there is also a down side to it, both at a personal level and on a more global level. Naturally, if users are not tracked they cannot enjoy a personalized viewing experience. While this may be fine for a small percentage of users, if too many people choose not to be tracked, this would have a global effect on the advertisement business.

In light of the above, we present a middle-ground approach in which users are able to enjoy a personalized viewing experience by receiving targeted ads and advertisers are able to know the amount of exposure that their ad received, while at the same time protecting the privacy of users. The way we cope with these seemingly contradictory requirements is by:

- Moving the processing of personal information

Reinman T. and Waisbard E..
PPiTTA - Preserving Privacy in TV Targeted Advertising.
DOI: 10.5220/0004076103270332
In *Proceedings of the International Conference on Security and Cryptography* (SECRYPT-2012), pages 327-332
ISBN: 978-989-8565-24-2

from the broadcaster[1] to the user (for example, to the set-top box of a TV viewer).

- Reporting back a statistically noisy reply that protects the individual, but still provides a good approximation for the advertisers.

Until recently targeted advertising was deployed mostly for web-based content, but now with the wide distribution of personal video recorders (PVRs), the set-top box is able to receive a variety of ads and to decide locally which ones to display. This localized ad insertion during a commercial break allows multiple TV viewers to watch different ads at the same time. While our approach is applicable to several settings, we focus on the TV + PVR case, as this is still the major content consumption platform.

As mentioned, the first step towards achieving user privacy with targeted advertising is to store and process ads on the user's device. The second step is to have multiple ads delivered to the user's device, each ad being selected based on a metadata analysis to determine the attributes of the optimal user presumed to be watching the particular program being viewed (for example, a program targeting mainly women should deliver ads related to women). The user's device would then find the best match between the user's profile and the ads delivered and display the ad that has the closest characteristics to the user's profile.

We later discuss how to parameterize the user's profile and to perform the matching with the delivered ads. If we were only interested in privately matching ads to users, this would be a final step. However, a crucial requirement is for the system to report back which ads were displayed. It was shown by Korolova (Korolova, 2010) that reporting which ads were viewed leaks a lot of private information. Thus, we need to find a way to provide this feedback to operators (which they can convey to advertisers) without leaking private information. A helpful observation is that advertisers are not interested in particular users, but rather in getting a global view of how many people actually watched an ad. By masking feedback reports with statistically noisy responses we can provide a good estimation of the global exposure that an ad received while maintaining user privacy.

2 RELATED WORK

The conflict between targeted advertising and consumer privacy arose in the context of web browsing. An advertiser's interest is to collect as much client information as possible, in order to display advertisements more efficiently. In contrast, the consumer is interested in preventing the exposure and exploration of private information by the advertiser[2].

Juels (Juels, 2001) discusses this conflict. He suggests solutions to reconcile these contrasting interests by using a *negotiant* – a client-side agent that mediates between the consumer and the advertiser. Basically, the idea is that the consumer profile is accessible to the negotiant, and not to the advertiser. Based on the consumer profile, the negotiant decides which advertisements to request from the advertiser – without disclosing the profile to the advertiser. In addition, if we would like to camouflage the ad category of a specific consumer from the advertiser, Jules suggests to use tools such as Mix Networks (introduced by Chaum (Chaum, 1981)), which assumes the honesty of (most of) the servers on the route.

Spangler et. al. (Spangler et al., 2003) present a way to profile households and TV viewers, based on their viewing patterns, in order to deliver targeted advertising. Their discussion focusses on data mining aspects, and barely refers to the privacy issues (they only recommend a privacy policy).

Toubiana et. al. (Toubiana et al., 2010) introduced Adnostic (http://crypto.stanford.edu/adnostic/), a system for targeting web advertising based on a user browsing behavior. Adnostic is a web-oriented online system that uses properties of web-browsing, such as cookies, to create a user behavioral profile in the user's browser.

The concept of differential privacy was introduced by Dwork et. al (Dwork et al., 2006). They gave a measure to the amount of privacy leakage. Loosely speaking, ε-differential privacy is designed to protect the privacy between neighboring databases which differ only in one row. As one can see in (Dwork, 2008) many papers have been written in this area, giving generic constructions. However, they address only one aspect of our problem. In our work, we take a more specific look and give a complete solution to the issue of targeted advertising, including user profiling and privacy preserving report back. Our construction is simple and efficient and can be implemented in real systems.

3 PROBLEM DEFINITION

The contradictory interests that need to be reconciled

[1] We interchangeably use the terms *broadcaster* and *operator*.

[2] Advertisers claim that there is also a consumer advantage to collecting client information – a typical consumer prefers relevant advertisements that target his needs and interests rather than those which are of no interest.

are the following: On the one hand, the advertiser needs to know information about the individual and the household so that advertising can be targeted. On the other hand, the viewer wants to protect the privacy of himself and his household. In addition, there might be regulatory prohibitions to collect and/or process private information outside the individual's premises. User *viewing habits* are a valuable source of information for ad targeting, as they tell the operator a great deal about the household. The advertiser would like to use this information to better-target the viewers. The problem is that this must be done in a way that does not violate the privacy of the viewer.

Generally speaking, targeted advertising requires user profiling, which can be done in one of two ways: 1. Using information that the user explicitly gives (e.g., gender, age, etc.). 2. Dynamic profiling that is derived from user's behavior, for example, TV viewing habits – which gives a very good indication about the user's interests.

A privacy-preserving targeted advertising solution must ensure that the advertiser knows how much exposure each ad receives. As mentioned before, the advertiser is not interested in what individual users viewed. We make use of this fact in our solution in order to give the advertiser a good estimation of the overall exposure of each ad, while hiding the individual's ad exposure pattern.

4 OUR SOLUTION

Our solution comprises three actions:

1. A method for dynamic user profiling
2. Determining which ad to present, according to the user's profile
3. Doing the above two actions on the user's device in a manner that preserves his privacy

The PPiTTA design enables targeting advertisement according to the user viewing habits – without violating his privacy. In addition, it preserves both individual privacy and global statistical accuracy. This is achieved by maintaining an approximated global view (by adding random noise) that does not reveal the choices of individual users.

4.1 Building the Profile

The first element of targeted advertising is the actual building of the profile. Basically, there are four methods to build the profile (all of them may be used jointly). The focus of this paper is the last method.

Profile Knowledge based on User Demography. When consumers subscribe to a broadcasting service, they are required to fill a form which includes information such as their age and number of people in the household. In addition, the specific package they purchase provides another indication about the household. There is also statistic information that can be used by the operator. For example, the consumer's neighborhood, the size (and worth) of the house, and the credit-card class may all be correlated with a socio-economic status.

Profile Knowledge based on the Users' Voluntary Disclosure. Advertisers, broadcasters, and survey companies may ask users to voluntarily (or for payment or other benefits) supply information about themselves. This information may include the level of income, consumer interests and history (for example, whether a user has bought a car or a house in the last year), hobbies, cultural outlook, and so on.

User Self-profiling. A user may ask for specific types of advertisements.

Behavioral-profiling based on User Viewing Habits. This method assumes a correlation between viewing habits and consumer habits and interests. Moreover, it assumes that this correlation is simple enough, such that valuable information about consumer interests can be deduced from viewing habits. The general idea is to analyze the viewing habits, and to profile the user accordingly. For example, the more sports shows a person watches, the more likely he would be interested in sports goods.

4.1.1 Profile based on Viewing Habits

A user can be profiled as a male or a female, as a teenager or an adult, and so on. A way to maintain a user profile is to have a list of attributes, such that for a given user, each attribute gets a score. We denote this list as the *attributes vector*. The attributes vector may have entries for demographic characteristics such as age, level of income, neighborhood, and entries for level of interest in different areas such as sport, fashion, and gadgets. The attributes vector is initialized to certain values, based on the package purchased by the user and known demographic information (region of residence, income level, education level, etc.), and is stored in the STB. Attributes whose values cannot be deduced from a-priori knowledge are initialized to some predefined initial values.

Whenever a user watches a TV program, the attributes vector is updated on the client device (e.g., STB) as follows. Initially, the content provider or the broadcaster gives each TV program a score for each of the attributes, meaning that each program has its own values for the attributes vector. For example, if a certain movie is known to have more appeal to teenagers, then the attributes vector for that content is marked with a higher score in its teenager-related entries. For certain content, some attributes will have no distinct categorization, and thus their score for those attribute will be zero.

When a user watches a TV program, the program's attributes vector is added to the user's attributes vector, and the attributes vector is then re-normalized (each attribute is normalized such that its score is within its valid range). By accumulating the score of these attributes, we get a high level of assurance about the viewer's profiling.

Let us see a more detailed example: John Smith purchased the Sports Gold Package with extra sports channels, including live events in HD quality. On his registration form, the following information has been disclosed: He is 30 years old and single, his income is in the top echelon, and he lives in the center of London. Consider an attributes vector with the following attributes (with the value range in parenthesis, and default value in square brackets – see Figure 1(a)): Age $(0-99)[0]$, Gender (M,F,B)[B][3], Number of people in household$(1-20)[4]$, Income$(0-10)[5]$, Sport$(0-10)[5]$, Fashion$(0-10)[5]$, Cars$(0-10)[5]$, Travel$(0-10)[5]$. With the current known information, the broadcaster could initialize John's attributes vector as follows (see Figure 1(b)): Age = 30, Gender = M, Number of people in household = 1, Income = 10, Sport = 8. Using the default values, the rest of the attributes can be initialized: Fashion = 5, Cars = 5, Travel = 5. (Alternatively, the broadcaster can base the initialization of these last attributes on heuristics of what a 30-year-old, rich, single sports fan would be interested in). Now John starts to watch TV and, as expected, he watches many sports program. On the other hand, he never watches fashion shows, and watches only a few cars-related and travel-related shows. His attributes vector is updated accordingly – see Figure 1(c). Then John meets a girl, Jill, and after a while she moves into his apartment. Jill loves fashion TV shows (so the fashion attribute score goes from 0 to 9), and John watches fewer sports programs than before (so the sport attribute score goes from 9 to 8). They also begin to plan a trip and start to watch travel programs, so the travel attribute score is increased accordingly – see Figure 1(d).

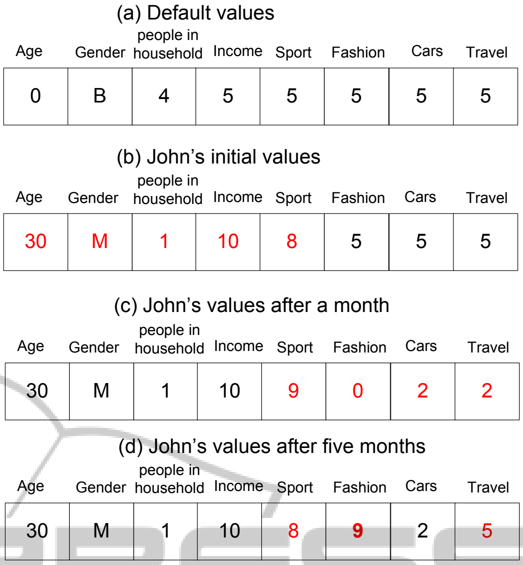

(a) Default values

Age	Gender	people in household	Income	Sport	Fashion	Cars	Travel
0	B	4	5	5	5	5	5

(b) John's initial values

Age	Gender	people in household	Income	Sport	Fashion	Cars	Travel
30	M	1	10	8	5	5	5

(c) John's values after a month

Age	Gender	people in household	Income	Sport	Fashion	Cars	Travel
30	M	1	10	9	0	2	2

(d) John's values after five months

Age	Gender	people in household	Income	Sport	Fashion	Cars	Travel
30	M	1	10	8	**9**	2	5

Figure 1: An Attributes Vector.

4.2 Deciding Which Ads to Display

The advertiser defines a target audience for each advertisement. This definition uses similar scores as the attributes mentioned in the previous section. For example, an advertisement of athletics is intended for males, who are interested in sport and the Sport attribute will get a high score. It is also possible to define attributes as "irrelevant" for a specific ad.

Initially, users are profiled and ads are categorized into profiles too, so only the matching still need to be done. To ensure that the user's privacy is preserved, all profile matching is performed in the user's device (STB). The STB stores a large collection of ads[4] of many kinds. According to the user's profile, the STB decides which ad to display and when. Each ad has its value (score) for each attribute, and the STB chooses the ad which is the most suited to the user's profile.[5] Ads are pushed to the STB ahead of time and are embedded into the viewed program (using an "ad-insertion" mechanism) in the commercial slots (or displayed in a banner or in the menus). As such, users (with different profiles) who watch the same program at the same time will get different advertisements, according to their profiles.

[3] Profiling is actually computed for the household, rather than for a single user so 'B' stands for 'both'. Research is being done to enable distinguishing between different users in the same household according to the time of day, the viewed content, and even the zapping speed and frequency. Machine learning techniques can be used to figure out who is likely to be the current user.

[4] Actually, the STB may use the same logic (of profiling and scoring) to select which ads to *store*, instead of storing many ads and select which of them to *display*.

[5] To be precise, suitability is not the only criteria. Usually, the contract between an advertiser and a broadcaster defines many conditions and limitations, e.g., how many times and how frequently each ad should be displayed. The STB has to maintain counters and other mechanisms in order to obey to these (and other) limitations.

4.3 Advanced Reporting

In today's systems, statistics of TV ad viewing is mostly based on "Panels"[6] or "Samples". These panels are designed to be statistically representative of the population. When an advertiser wants to know how many people with a given demographic were exposed to a given ad, and how many times on average each individual of the target population saw the ad ("spread" and "reach"), they take the numbers for the Panel and extrapolate them to the population at large.

We propose an alternative approach that collects data from the entire viewer base and use privacy protected reporting to provide better granular spread and reach data.

4.4 Usage Flow Example

In this section, we demonstrate a use-case, in which the advertiser is charged according to the audience who watched his ad. We start with the pre-known static data, that is not related to viewing habits, and is usually collected by the operator in the registration/purchasing process and deduced from a publicly available data. Such data includes the wealth of people in the neighborhood, the percentage of families that own their home, etc. We start by setting the values of these attributes accordingly (e.g. setting the "Own home" to a score of 90% in a rich neighborhood).

Next, user viewing habits effect the user's attributes vector and this is used to deduce a user profile. For example, when a particular show is viewed regularly (e.g., MTV hits), there is an 80% chance that there is a teenager in the home. Every time this show is viewed, the "Teenager in home" score goes up a bit. This is considered to be private data that viewers do not wish to share. Naturally, combining both static and dynamic data allows a more accurate prediction (combined also with estimations from Panel data).

The price that advertisers pay is relative to the exposure it gets. The measuring unit is CPM (Cost Per Mil, i.e., cost per thousand). The more targeted the ad is, the more an advertiser is willing to pay. For example: 5$ for sports fans, 6$ for male teens, 7$ for male teens that are sports fans, 10$ for male teens that are sports fans and that have affluent parents.

Using our technique, the CPM will be set according to the viewer's profile. By maintaining an ongoing count of how many times each ad was presented to a particular user, we would be able to plot a detailed chart. This chart describes the number of times this ad was presented, how many viewers were exposed to it at least once, and how many saw it many times (e.g., more than ten times).

4.5 Privacy Preserving Impression Report

Another challenge is to report back statistical usage while preserving both individual privacy and global statistical accuracy. These are two somewhat contradictory tasks: the first being to provide an accurate report of users' consumption and the second being to preserve the privacy of each user. Here we want to achieve differential privacy, meaning that the global reply is not affected by more than ε fraction by the inclusion or exclusion of any specific user.

We construct two schemes that achieve this privacy. Each scheme is designed to be robust against a different type of adversary that may have different a-priori knowledge about user behavior.

4.5.1 Individual Ad Mask

The first scheme hides the exact number of impressions of each ad A_i that the user j consumes by adding a random value r_i^j to its real impression counter c_i^j and reports back $\bar{c_i}^j = c_i^j + r_i^j$. In our construction, we recommend using a Gaussian distribution.

Computing the sum of the impressions for A_i is:

$$\sum_j \bar{c_i}^j = \sum_j c_i^j + \sum_j r_i^j$$

Note that $\sum_j c_i^j$ is the exact sum of the c_i^j, and $\sum_j r_i^j$ is a random combination that is treated as a random noise with its mean equal to zero.

4.5.2 Accumulated User Impression Counter

The second scheme gives better protection for users who may have more extreme impression counters. In this scheme, each viewer reports back a value that is the weighted sum of all his counters. Now, even a high value counter is well masked by the other counters of that individual. The scheme works as follows: Let n denote the number of users, k - the number of different ads, and $c_1^j, \dots, c_k^j$ be the counters of the impressions for a user j.

1. $\forall j \in \{1 \dots n\}$, user j would pick a random vector $v^j = v_1^j, \dots, v_k^j$ where each $v_i^j \in \{-1, 1\}$
2. User j would compute $s^j = \sum_i v_i^j \cdot c_i^j$
3. User j would report $< s^j, v^j >$

[6] A "panel" is a subgroup of the population who agree to have their data collected and analyzed.

The operator, after collecting all the reported counters $< s^1, v^1 >, \ldots, < s^n, v^n >$, is able to compute an estimated impression counter as follows:

$$\bar{s_i} = \sum_j v_i^j s^j = \sum_j (v_i^j \sum_t v_t^j \cdot c_t^j) = \sum_j \sum_t (v_i^j v_t^j c_t^j) =$$

$$= \sum_j \left(\sum_{t=i} (v_i^j v_t^j c_t^j) + \sum_{t \neq i} (v_i^j v_t^j c_t^j) \right) = \sum_j c_t^j + random$$

Since v_i^j are chosen at random, for $t \neq i$ the summation acts as a random value; thus we get a noisy sum of c_i^j.

5 SECURITY ANALYSIS

In this section, we analyze the security of our scheme and emphasize the advantages in the context of the two adversarial models. We start with analyzing the first scheme, which hides the exact number of impressions for each ad. Clearly, this scheme achieves privacy, since adding a report of an additional user to the server's database does not change the overall distribution of impressions in the database.

By adding Gaussian noise with a distribution (μ, σ^2) we get that

$$\sum_j \bar{c_i}^j \sim (\mu, (\frac{\sigma}{\sqrt{n}})^2)$$

The larger σ is, the better privacy protection we get for the user. The smaller $\frac{\sigma}{\sqrt{n}}$ is, the more accurate is the estimation of the total number of impressions. So by carefully choosing σ, one can set the tradeoff between accuracy of the result and the level of privacy.

However, there is one more thing we need to address – that is the potential privacy leakage over time. Consider the following attack: The operator keeps a log of all the reports he gets from a single user. If he notices that for a specific category, such as sports, the user's report is often above the average impression count, then he can deduce that the user is a sport fan. The way to cope with this is by periodically deleting the logs after they are processed. We claim that it is reasonable to perform this because we do not consider the operator as malicious, but as semi-honest. If the operator was malicious, then could simply have the set-top box report back the exact values (the operator controls the set-top box software). While the operator is assumed to act in good faith and to follow the privacy regulations, he does not want to retain private information any longer than necessary (and potentially have it exposed to insiders).

Our second scheme avoids possibility of long-term learning even if the operator does not delete the logs. However, there is a different potential weakness in this approach. Consider the following attack: Assuming the number of possible ads is small and that an adversary has auxiliary information about a particular user (e.g., via the package he has purchased), then the adversary can make a good estimation about the distribution of the different impressions for that user. Here again we claim that the operator is semi-honest and would not store such information about the user. Therefore, it is unlikely that insiders would have such a-priori knowledge about individual users, and outsiders with such a potential knowledge would not have access to the logs.

6 CONCLUSIONS

In this paper, we have presented the first practical scheme for achieving targeted advertising in TV systems, while preserving the user privacy. We showed how to build a household profile, how the set-top box decides accordingly which ad is the most appropriate to display, and how to report back the impressions to the operator in a privacy preserving manner.

REFERENCES

Chaum, D. (1981). Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–88.

Dwork, C. (2008). Differential privacy: A survey of results. In Agrawal, M., Du, D.-Z., Duan, Z., and Li, A., editors, *TAMC*, volume 4978 of *Lecture Notes in Computer Science*, pages 1–19. Springer.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In Halevi, S. and Rabin, T., editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer.

Juels, A. (2001). Targeted Advertising ... And Privacy Too. In Naccache, D., editor, *CT-RSA*, volume 2020 of *Lecture Notes in Computer Science*, pages 408–424. Springer.

Korolova, A. (2010). Privacy violations using microtargeted ads: A case study. In Fan, W., Hsu, W., Webb, G. I., Liu, B., Zhang, C., Gunopulos, D., and Wu, X., editors, *ICDM Workshops*, pages 474–482. IEEE Computer Society.

Spangler, W. E., Gal-Or, M., and May, J. H. (2003). Using Data Mining to Profile TV Viewers. *Commun. ACM*, 46(12):66–72.

Toubiana, V., Narayanan, A., Boneh, D., Nissenbaum, H., and Barocas, S. (2010). Adnostic: Privacy Preserving Targeted Advertising. In *NDSS*. The Internet Society.