Adaptive User Profiles in Pervasive Advertising Environments

Florian Alt[1], Moritz Balz[2], Stefanie Kristes[3], Alireza Sahami Shirazi[1],
Julian Mennenöh[3], Albrecht Schmidt[1], Hendrik Schröder[3], Michael Goedicke[2]

[1]Pervasive Computing and User Interface Engineering Group
[2]Specification of Software Systems
[3]Marketing & Trade
University of Duisburg-Essen
Schuetzenbahn 70, 45117 Essen
{florian.alt, alireza.sahami, moritz.balz, stefanie.kristes,
julian.mennenoeh, albrecht.schmidt}@uni-due.de

Abstract: Nowadays modern advertising environments try to provide more efficient ads by targeting costumers based on their interests. Various approaches exist today as to how information about the users' interests can be gathered. Users can deliberately and explicitly provide this information or user's shopping behaviors can be analyzed implicitly. We implemented an advertising platform to simulate an advertising environment and present *adaptive profiles*, which let users setup profiles, based on a self-assessment, and enhance those profiles with information about their real shopping behavior as well as about their activity intensity. Additionally we explain how pervasive technologies such as Bluetooth can be used to create a profile anonymously and unobtrusively.

1   Introduction

In 2006, only 41% of the expenditure on advertising actually produced sale [5]. Modern advertising environments such as online stores or networks of public displays, try to increase this ratio by not distributing or presenting advertisements to users randomly but by targeting contents towards users based on their interest. Therefore, marketers need to know their customers in order to decide which products to present where and when [6].

A common approach is that advertisers define target groups for their ads, similar to companies defining target groups for their products in general. When it comes to assigning users to specific target groups, two approaches prevail today as to how information about the users' interests and shopping behavior can be gathered and later be used to enhance the exposure of advertisements. On one hand, users can deliberately provide these data; on the other hand, data can be gathered by analyzing the users' shopping behavior. Based on the profiles target groups are created, which might then be selected for a specific campaign. Though target groups are an easy and effective way of tailoring ads towards the interests of a user group, we think that when abstracting from a user to a target group, important information gets lost that might allow for a more precise and successful targeting.

In this paper we present a platform we implemented to simulate an advertising environment in our lab. Our current setup allows for adapting advertisements towards users' interests and also lets users interact with public displays we use to show the advertisements. As a contribution we present three techniques which aims at tackling shortcomings of current approaches.

First, we present the approach of *adaptive user profiles*. It allows users to setup a profile based on a self-assessment. This profile is then automatically enhanced with data about the "real" shopping behavior. In this way, on one hand the users' subjective interests are taken into account. This is especially interesting for advertisers since it reflects the users intention and provides a strong indication whether or not a user is interested in a specific (type of) product. On the other hand, an implicit observation of the user might reveal behavior or intentions a user is not aware of or simply cannot remember or articulate. Hence a more detailed image of the user can be drawn and advertisements be targeted more precisely as if simply assigning users to target groups.

Second, we explain how based on this sophisticated method of generating user profiles we can compare a user's activity against other users by calculating a *user activity intensity*. We use this measure when calculating the adaptive user profile and discuss its value for an advertiser.

Third, we show how pervasive technologies such as Bluetooth can be used to *create profiles in an anonymous and unobtrusive way*. Linking profile information to the MAC address of a Bluetooth enabled device, such as a mobile phone, allows for identifying a user without the need to store any personal data such as name, gender, or address. Further, a user may easily discard from using the system by switching off the Bluetooth functionality of his phone.

The paper is structured as follows: first we present the approach of adaptive user profiles, show how such profiles can be generated and critically reflect on privacy. Second, we come up with a technical concept and show how we implemented our system. Finally, we present related research and discuss future work.

2 Adaptive User Profiles

In this chapter we first provide an introduction to profiling and target groups. Second, we present an approach for creating *adaptive user profiles*. It allows for generating profiles based on user-generated information and data derived from the users' interaction with an advertisement system. Then we explain how those profiles can be used to match ads towards users. Finally we consider the aspect of privacy and show how anonymous profiles can be used to avoid users' concerns regarding the abuse of personal data.

2.1 Introduction to Profiling and Target Groups

In order more precisely target content, information about the customer is required. Such information are gathered and stored based on profiles. In general, profiles may be created in two different ways. First, users can explicitly share this information with the advertiser. Many online stores explicitly ask users about their areas of interest when setting up their personal profile or when signing up for shopper loyalty programs. Sec-

ond, data can be derived from the user's shopping behavior. While in online shops this can be done implicitly without the requirement for any user interaction, e.g., by analyzing a user's click-stream, the user's cooperation is required in stores, e.g., by presenting their shopping loyalty card upon checkout.

Based on the profiles it is possible to form target groups by applying cluster analyses on the overall data set. Advertisers might then select one or more pre-defined target-groups for their campaigns.

2.2 Creating Profiles

Our approach generates profiles, which consist of ratings on a scale from 1 to 10 in several categories. Those categories depend on the advertising environment and may be adjusted to the needs of the provider, based on an analysis of the consumers or a market analysis. In our lab setup, we specified the categories for a local-scale advertising environment including local stores, events, etc. In total, ratings for 45 categories reflecting the users interest may be calculated. They include for example music, events, and places and are used to define different types of profiles: the *consumer profile*, *campaign profiles,* and the *target group profile* for a specific campaign.

Consumer Profiles

The consumer profile reflects the user's areas of interest and consists of two components. The first component is a self-assessment, allowing users for specifying their interest in the given categories ranging from "not interested" (1) to "very interested" (10). By providing this explicit profile to the system, a user can benefit in a way such that ads he receives are more appealing to him than random ones since they match his interests.

The second component tries to take into account the actual shopping behavior. Due to the fact that every user interaction with the system is captured, it is now possible to enrich the user's self-assessment using his behavior. The system currently tracks the following types of interaction for each category per user: *Point of Sale Visits* (A), *Event/Ad Watching on a Display* (B), *Rating of an Event* (C), and *Redemption of Coupons* (D). This set of interactions is not conclusive and further types of interaction may be considered.

Assuming that a user's visit of a sport store (A) signalizes his interest for the category "sport", this store visit should positively affect the user's profile in the assigned categories. Similarly this is true for event/ad watching (B) and the selection of coupons (D) offered by the system as a motivation to participate. We have not yet integrated the event rating (C) into the system, because we are unsure if or how the rating can be translated into effects on users' interests. However, the impact of a certain type of interaction depends on measures such as the resulting revenue. Since we did not yet deploy and evaluate the system in the real world and hence do not have any data, we cannot provide evidence that our model correctly reflects the influence of the different types of interactions on the consumer profile. For example, can we really derive a higher interest from more frequent store visits or can we conclude that somebody's interest is lower because he missed to use his coupon on various occasions?

The consumer profile used for targeting the advertisements is calculated in three steps:

Calculating Individual Activity: For each user data is available showing, which action took place in which category. Our system aggregates the different kinds of actions to get *one* value per user per category indicating how active a user behaves with regard to a certain category. An aggregation can be displayed as a linear combination:

$$Individual\_Activity = w_1A + w_2B + w_3C + w_4D$$

with w = weights and A,B,C,D = number of actions

The weight parameter is required since a scaling is necessary depending on the impact of the action. For example, the overall amount of the redemption of coupons will be on average much lower than the amount of point-of-sale visits. However, this does not indicate that coupons are less important. Hence, the weight parameter can be used to "equalize" the different distributions in relation to their means.

Once data are available, such as a positive consumer reaction to an ad or the revenue generated by the user in this category, the weights could be estimated by using a multiple stepwise linear regression analysis. Subsequently, the weights would indicate the "importance" in relation to the influence of the action type on business relevant target measure such as the revenue. In the current state, we still have to integrate a measure to track the real impact of the advertisement, e.g., the actual sales possibly triggered by the ad.

Calculating Individual Activity Intensity: To estimate the overall user activity per category we suggest the measure *individual activity intensity*. It reflects how much the individual activity differs from the overall mean activity of all users.

For example: a value of 1.2 (0.8) would indicate that the activity lies 20% above (below) average within a certain time period. Users having a high individual activity intensity might be of special value to the advertisers since it can be assumed that their probability to generate revenue is highest among those users.

Adjusting the Profile: Since for each user the self-assessment of interests is available, the consumer profile can now be enriched using the activity intensity. An obvious approach would be to linearly increase or decrease the user's self-assessment based on whether his activity intensity is above or below average. However, this might lead to ratings beyond the scale. We use an adjusted normal distribution to assure that the value always stays within the boundaries of the scale. The profile can be calculated as follows:

$Consumer\_Profile =$
$\quad Profile_{self\_assessment}$
$\quad + f(Profile_{self\_assessment} \cdot activity\_intensity)$
$\quad \cdot ((Profile_{self\_assessment} \cdot activity\_intensity) - Profile_{self\_assessment})$

with f(x) = adjusted normal distribution

The formula uses the self-assessment profile as an input and adjusts the value based on the user's activity intensity. The normalization function f controls how quickly or slowly the profile value changes. We plan to use the values from the campaign profile to give advertisers an influence as to how strongly their campaigns alter the users' profiles.

Campaign Profiles

Campaign profiles consist, similar to consumer profiles, of one value in each category, which can be determined by the advertiser upon creating the campaign. Its use is to characterize advertisements, coupons, and points of sale belonging to the campaign. This characterization is required to determine which category values of a user are to be adapted through a certain interaction. For example, the campaign of a sport retailer will have high values in the sport category. A consumer's visit of a point of sale, which belongs to the campaign or the retailer, will subsequently have an effect on the rating of his category sport. Since the campaign has a value zero in cultural events, a point of sale visit will not alter this category.
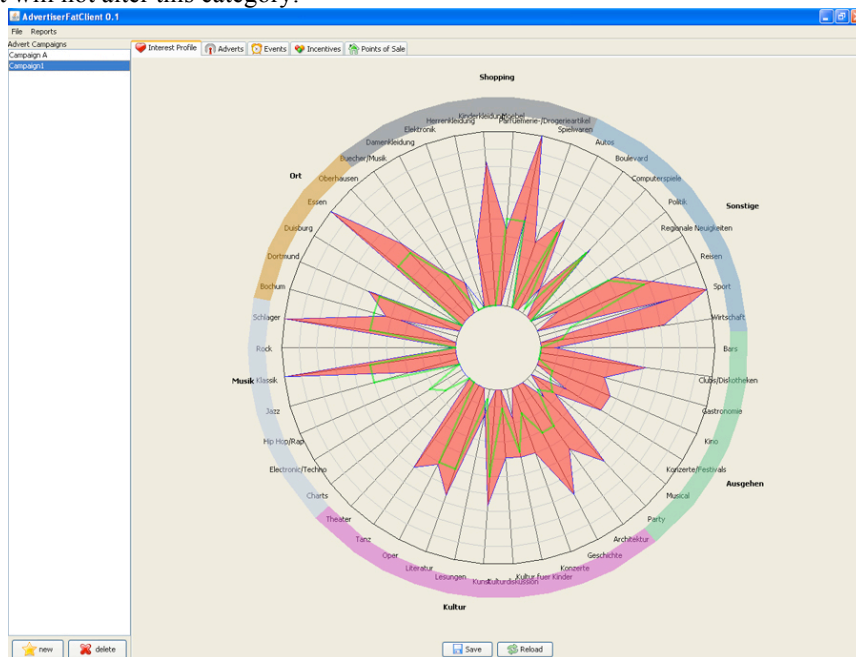

Figure 1: Advertiser's client: Advertiser's can specify their campaign's profile (green line) defining the influence of the campaign on a user's profile, as well as the target group profile (red area) defining which user group the campaign is targeted to.

Target Group Profiles

While the consumer profiles are calculated as a blend of self-assessment and actual user interaction with the system, advertisers can determine target group profiles using minimum and maximum values for each category. In addition, advertisers need to define a campaign profile for each of their campaigns, which, like consumer profiles, only hold one value for each category. This target group profile is then used to match the campaigns with the user's interests. Since advertisers can select from a range of values, they may define precisely which users they want to target the ads to and which not. For example, if a sports retailer wants to target only consumers with a high interest in sport,

they might define a campaign with a target group profile ranging from values 8 to 10. **Figure 1** depicts our advertiser's client, which allows for specifying the campaigns. The red area shows the group of consumers the campaign should be targeted to (e.g., users with a value in the sports category between 8 and 10). The green line indicates how a user's profile is affected upon interacting with content related to this campaign.

An interesting idea from the marketing point of view is to identify users that stated a high interest in a certain category during the self-assessment of their interests, but do not show any activity with the system in this category. By addressing those users with an advertisement in that category, chances might be higher to trigger activity since general interest can be assumed. This case shows that having both a self-assessed user profile and an adapted user profile is very valuable for advertisers.

2.3 Targeting Content towards the Users

Based on the target group, ads to be displayed are chosen. We distinguish two cases when it comes to matching ads and users. If only a limited number of users are in the display's vicinity, ads are selected by showing the best fitting ads for only one single user taking turns in order of arrival. The main reason is that finding a best match for a small group of users might be difficult if interests among them differ fundamentally. Hence, a more promising strategy is to perform a "safe" match on the single users. With an increasing number of users at the display, we use matching algorithms to get an over-all sufficient fit. Both cases require a measure indicating how well an ad fits a (group of) users. The distance measure should meet the following requirements:

- If the user's category value lies within the defined borders of the target group profile, the distance should be 0. Thus, the distance is measured between the individual value and the boundary, if the value lies beyond. Else, the distance is 0.
- It should be considered that a larger distance might indicate an aversion towards a certain category. So, it is suitable to choose a measure using squared values.

Based on those requirements, we suggest using the squared City-Block-Distance, which sums up the squared distances for each dimension. This metric takes into account that larger distances might be disproportionally worse than smaller distances. Hence users' objections towards specific categories are well considered and the probability of targeting advertisements towards users he does not like is being minimized.

2.4 Privacy

When it comes to storing and assessing information about a user's shopping behavior, this often raises privacy concerns. Consumers are afraid of data being passed on to other companies or advertisers hence resulting in an ever-growing amount of advertisements users are confronted with. However, the success of shopper loyalty programs shows, that most of the users are nevertheless willing to give away their information in return for incentives [2][8], even though it is not entirely obvious who has access to the data.

With our approach we show that it is possible to target consumers without the need to store personal data such as name, age, gender, or address by linking a profile to a Bluetooth MAC address only. Thus, the user stays anonymous, and may at the same time

deliberately decide whether or not he wants to participate in the advertising environment. This model, however, makes it essential for the success of the system to convince the user of its advantages, such as advertisements, which are personally interesting for him, and make him aware of the fact that he is virtually in control of the system at any time.

3 Technical Concept

In order to apply the profiling approach described above, we need a system that is capable of recognizing users and relating their interests to advertisements. We will first outline the requirements and general functionality of such a system; our concrete implementation will be described afterwards.

3.1 Requirements and Approach

The most important feature of this system is the management of users and their profiles. To be as pervasive as possible, users should not be represented explicitly with personal information. Instead, they should be identified implicitly whenever they connect to the system; for this purpose the identification mechanism must define a relation to the connection technologies. When a user is identified, the system can relate any data regarding profiles, interests, and advertisements to this abstract profile.

Since the system will act pervasively, the connection mechanisms play an important role, since they allow to identify and track users; an example is the MAC address of a Bluetooth device a user can utilize to connect to the system. However, we also want the system to be as flexible as possible. This means that a wide range of connection mechanisms should be supported. For each deployment of the system, a set of appropriate technologies can be selected that matches the users' target group. To illustrate this, the following exemplary scenarios may be considered:

- The system is deployed at a traffic junction. Waiting persons can connect to it with Bluetooth-enabled mobile phones. The user range is thus limited to persons carrying such phones and being able to operate Bluetooth software, but advertisers can assume to interact with people having a technological background and can afford advanced hardware. Since no other assumptions of the user base can be made, Bluetooth devices are the lowest common denominator. Users are identified by the unique Bluetooth MAC address and can thus be recognized each time they pass by.
- The system is deployed at a university. Students at this university are used to pay in cafeterias with a student ID card that is equipped with an RFID chip and a barcode denoting the matriculation number. The system will read either of them; the RFID chip can even be used to detect users in the vicinity of a terminal. The user base is thus limited to students of this university; however, the students can connect to the system by just using the ID cards.

The requirements are thus conflicting: Connection mechanisms must be considered to identify users, but can be very different and should not affect the architecture or limit the functionality of the system. This means that two software components are required to handle connections: First, an abstraction layer is required that decouples connections

from the actual system by providing user identification at the same time; second, we must allow users to change their preferences and match different login credentials if more than one connection mechanism is offered.

The objective of the system is to employ this basic user management to offer matching advertisements. This is done by the profiling approach defined above. The system must thus handle the related data: For all users and ads, a set of *categories* must exist. Whenever a user takes an action that is of interest to the profiling algorithm, the profile values related to the categories are adjusted. While this is sufficient for basic functionality, we also want input from users. Hence users should be able to manage their own profile and give a self-assessment, which is used as an input to the profiling algorithm.

When the system fulfills these requirements, it provides all information to advertisers for target-group-specific ads. An appropriate *ad management* allows for supplying the information about marketing campaigns. These campaigns are defined by relating them to the categories to define user interests they will match. Each campaign consists of a set of media that can be displayed to the user by display clients that connect to the system.

3.2 Interaction types

As defined in the requirements, the user interaction must be considered carefully. Hence we defined the four types of user interaction with our system. We analyze these actions and use them to enhance the users' profiles.

- *Point of Sales visits*: Bluetooth scanners attached to the system are capable of sensing mobile devices in their vicinity. Hence it is possible to gather information about shops a user visited trying to find out his preferred brands.
- *Event / Ad Watching on a Display*: Users may want to read information on an upcoming event, such as a rock concert. This can be done explicitly at displays by using a mobile phone application.
- *Event Signup and Rating*: Once a user reads information about a potentially interesting event he may sign up for this event. Though it is not (yet) possible to buy tickets for events, feedback from the users in the form of ratings or comments may be gathered asking them later to rate the event later.
- *Incentive Choice (Display Chase)*: As a motivation for users to install client software on the mobile phone, a so-called display chase is offered. Users need to walk along multiple displays and answer region-related questions. Upon successfully answering the questions, the system offers them three coupons users can choose from. Coupons can be redeemed in nearby stores.

4   Implementation

In the following chapter we describe the different components of the overall system architecture is depicted in Figure 2. A *central server component* is used to persistently store advertisements, events, visitor profiles, and the users' interactions with the system. Based on this information the system can calculate a user's profile as described above. This profile is matched against advertising campaigns defined by the advertisers via a web-based client. The content can be outputted to any client, e.g., a public display. The
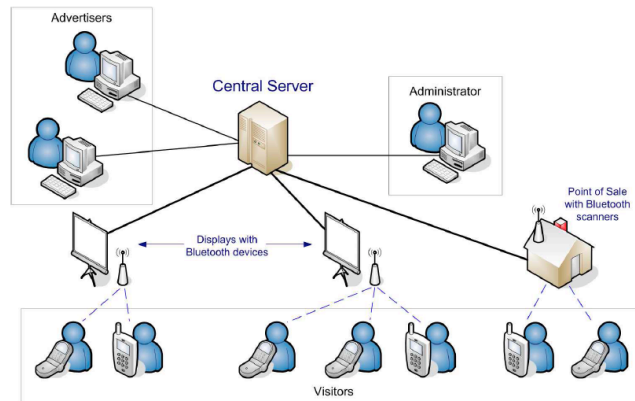
Figure 2: System architecture

advertiser's client allows advertisers to specify campaigns and select user groups to target ads to. This is done using the UI depicted in **Figure 1**.

The campaigns generated by the advertisers are presented to the consumers via public displays. Each public display is equipped with a Bluetooth scanner hence the system can "identify" users in front of the display and select content according to the profile related to the scanned Bluetooth MAC address. We implemented a mobile phone client in Java ME, which can be run on most Symbian phones. The mobile client can be downloaded via Bluetooth at the public displays. We use this application for two purposes: first, users can create their self-assessed personal profile. Second, the user can remotely control the display, e.g. browsing for information on a specific event, by using his mobile phone. We use Bluetooth scanners also in so-called points of sale in order to detect, which (areas of a) store a user visited. This reveals important information on the users' interests.

Finally, an *administration client* provides means to outputting the gathered data in order to create reports. These can be used by advertisers to evaluate the success of campaigns as well as by the providers of the system for accounting purposes.

5   Related Work

Advertising is not a new research domain in pervasive computing and has already been widely explored in recent years. Rangathan et al. showed how pervasive computing technologies could be used to deliver relevant ads in suitable ways to selected consumers [6]. Sharifi et al. [7] showed how public electronic displays can improve the efficiency of advertising systems if the display is aware of the identity and the interests of the audience. [4] focuses on how contents for public displays can be selected based on the audience's preferences and areas of interest. The ad selection within their system BluScreen tries to maximize the exposure of as many ads as possible to as wide an audience as possible without any prior knowledge of the audience and need of any user action.

Further research in this domain has focused on how available technologies can be used to provide more suitable and relevant information based on profiles. Bilchev introduced the concept of a distributed user profile constructed based on different profiles that can be used for personalized online advertising [1]. Stablere et al [8] presented a platform using information extracted from a Facebook profile to target personalised ads.

Last, user interaction with pervasive environments is an important aspect, which can be considered to generate profiles; however it heavily depends on the media. Whereas in

the WWW interactions can be easily assessed through analyzing a user's click-stream, tracking interactions with public displays requires additional sensing techniques. [7] explains how Bluetooth can be used to identify users in the vicinity of a public display.

6　Conclusion and Future Work

We introduced a pervasive advertising environment, which tries to tackle several issues of current deployments. First, we looked at how both the users' personal interests as well as their shopping behavior when interacting with the system can be taken into account in order to precisely target advertisements to them. Therefore we introduced the concept of adaptive user profiles. We showed how such a profile could be generated based on a user's self-assessment, the user's individual activity, and the activity intensity compared to other users of the system. Third, we proposed a system setup allowing for anonymous and unobtrusive profiling based on Bluetooth scanners. Finally we explained a technical concept based on which we implemented our pervasive advertising system.

Our advertising environment has so far been deployed in a lab setting only and hence not been evaluated with real-world data. As a future work we plan to setup the system on the campus of the University of Duisburg-Essen and run a large-scale user study in order to evaluate the different aspects of the systems. Those aspects mainly include (1) accuracy and robustness of the profiling algorithm, (2) performance and accuracy of the matching algorithm, (3) acceptance among user as well as potential user concerns, and (4) a usability evaluation of the different types of interactions supported by the system.

References

1. G. Bilchev, D. Marston, Personalised Advertising – Exploiting the Distributed User Profile, BT Technology Journal 21 (1) (2003) 84-90.
2. M. H. Bosworth. Loyalty Cards: Reward or Threat? ConsumerAffairs.com 2005 April 3, 2009. [http://www.consumeraffairs.com/news04/2005/loyalty_cards.html]
3. D. Kern, M. Harding, O. Storz, N. Davis and A. Schmidt. 2008. Shaping how advertisers see me: user views on implicit and explicit profile capture. In: CHI '08 Extended Abstracts on Human Factors in Computing Systems. ACM, New York, NY, 3363-3368.
4. M. Karam, T. Payne and E. David. Evaluating BluScreen: Usability for Intelligent Pervasive Displays. In Proceedings of ICPCA 2007.
5. L. Marsland: How much advertisement actually works? SAMRA Convention, bizcommunity.com, April 15th 2006. [http://www.bizcommunity.com/Article/196/119/9593.html]
6. A. Ranganathan and R. H. Campbell. Advertising in a pervasive computing environment. In: Proceedings of WMC '02, pp. 10–14, New York, USA, 2002.
7. M. Sharifi, T. Payne, and E. David. Public Display Advertising Based on Bluetooth Device Presence. In Mobile Interaction with the Real World (MIRW 2006), September 2006.
8. B. Spethmann. Shoppers Like Grocery Loyalty Cards Despite Privacy Worries. Promo Magazine 2005 [http://promomagazine.com/incentives/shoppers_loyalty_cards_020205/]
9. M. Stabeler, R. Shannon, and A. Quigley. Profiling and targeting opportunities in pervasive advertising. In Pervasive 2009 workshop on Pervasive Advertising, Nara, Japan, 2009.