User Profiling for Generating Bids in Digital Signage Advertising Auctions

Jörg Müller and Antonio Krüger

Institute for Geoinformatics, University of Münster, Münster, Germany,
{joerg.mueller, antonio.krueger}@uni-muenster.de

Abstract. In this paper we present a simple approach to adapt advertisements on digital signage to the interests of the audience. We use auctions to sell the advertising space to the highest bidding advertisement and decision theory to determine for each advertisement how much to bid. Within this framework, we use methods from content based recommender systems to adapt the bid to the interests of the audience. Each advertisement has a set of keywords, and the history of all advertisements a user was interested in is kept. We propose to use a naive Bayes classifier to estimate the probability that a user is interested in a certain advertisement given the keywords and users history. We propose interaction mechanisms to generate feedback to train the estimation functions and use the m-estimate to deal with the cold start problem.

1 Introduction

Imagine you walk around with two friends in your favorite shopping mall. You just pass a digital sign where two strangers are standing in front of it. As you notice that the display changes its content to adapt to your interests, you stop to have a glimpse on it. You see an advertisement presenting a discount for a DVD you are interested in, and take out your mobile phone to establish a connection to the display and copy it to your phone. While your mobile phone then navigates you to the place where you can buy the DVD, you wonder how the digital signs can always present you advertisements you are really interested in. In this paper we focus on how to select those advertisements you are interested in based on your interaction history. Main contributions are

- The application of the naive Bayes classifier to this problem.
- The presentation of feedback mechanisms to train the estimation functions.
- The application of the m-estimate to solve the cold start problem.

2 Related Work

The concept of using auctions to select advertisements shown on digital signage was introduced in [5]. The audience is sensed with a Bluetooth sensor, and advertisements are preferentially shown to those users that have not seen them yet. A bidding strategy for digital signage advertising auctions that uses the expected utility given a certain context to determine the optimal bid is presented in [3]. The application of the naive Bayesian classifier for recommender systems was proposed in [6]. In that work, websites are categorized into "interesting" and "not interesting" categories. The system is trained by manual feedback from the user and employs the naive Bayes classifier to automatically classify webpages based on features extracted from the content. An overview of recommender systems is provided in [1]. A general overview of work on situated public displays can be found in [4].

3 Auctions for Digital Signage Advertising

In this paper we concentrate on advertisements that entice the user to take a specific action, eg. buy a certain product. This action has a certain utility $U(act)$ for the advertiser, which

Fig. 1. The test environment for this study. The displays are installed in an university setting and show advertisements for talks, lectures and seminars.

is why the advertiser advertises at all. Because advertising space on displays is limited, for each situation the advertisements that are shown must be selected. We sell advertising space using a generalized second prize auction [2]. Each advertisement is represented by an agent, which is provided context data to generate its bid. For this paper, the context taken into account is the number of faces in front of the display, as detected with a camera, and the Bluetooth IDs of the audiences mobile phones, as detected with a Bluetooth sensor. Each advertisement should then bid its individual utility of being shown given the context data.

4 Applying User Profiles to Estimate the Optimal Bid

A simple model of the expected utility of showing an advertisement $EU(ad)$ is the utility that the user takes the advertised action $U(act)$ times the probability $P(act)$ that he does so. For multiple users in the audience, the individual utilities can simply be summed. Thus

$$EU(ad) = \sum_{i \in viewers} P(act_i | Context_i, ad) * U(act_i)$$

To be able to estimate whether a certain user is interested in an advertisement, we follow the approach proposed in [6]. Each advertisement is represented by a number of features f_i, in our case we use the submitting institution and a number of manually entered keywords. It would be helpful, although not necessary, if keywords are semantically unambiguous, that is, different advertisers use the same keywords for the same purposes. Then we take the naive Bayes assumption that the features are conditionally independent given act and ad, and apply Bayes' rule to achieve

$$P(act_i | f_1 \dots f_n, ad) = P(act_i | ad) \frac{\prod_{j \in feat} P(f_j | ad, act_i)}{\prod_{j \in feat} P(f_j | ad)}$$

$P(f_j|ad)$ is the probability that a new advertisement has the feature f_j and can simply be estimated from the history of advertisements. $P(act_i|ad)$ is the probability that the user acts upon any given advertisement. It can be estimated from the percentage of advertisements the user has seen up to now and those he acted upon. $P(f_j|ad, act_i)$ is the probability

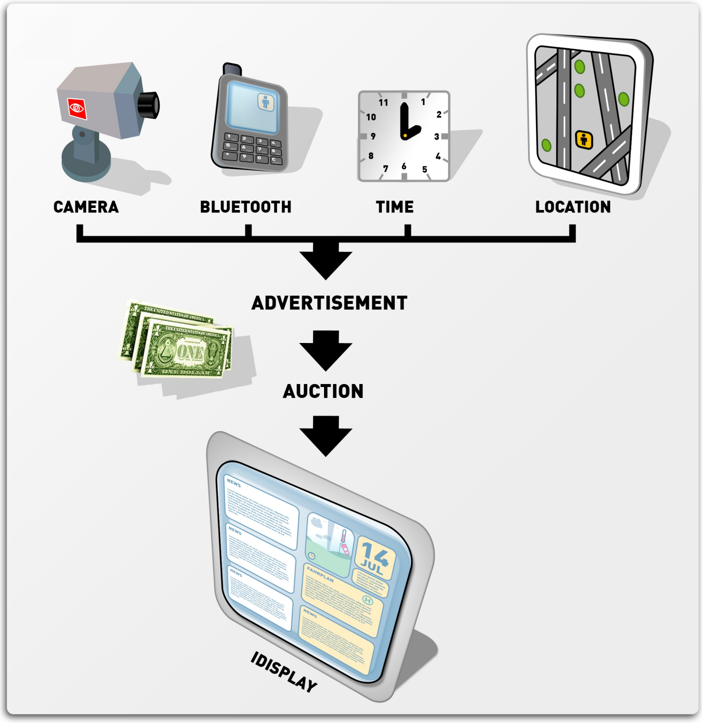

Fig. 2. The auction process used in our approach. At the start of each advertising cycle, context information is gathered. The number of users looking at the display, their identity, and the current time and location are provided to the advertisement. Advertisements take this context information and estimate the probability that the users will act. From this estimation they compute the total expected utility of being shown and generate their bids for the upcoming advertising cycle. In a generalized second price auction, the available advertising slots are then sold to the highest bidders. For the duration of the advertising cycle, the selected advertisements are shown on the display.

that an advertisement that the user acted upon has the feature f_j. This can be estimated from the past advertisements the user acted upon. To be able to estimate both of the last probabilities, a feedback mechanism is necessary to determine which advertisements the user really acted upon. We propose two approaches to implementing such a feedback mechanism. First, interaction mechanisms, for example via the user's mobile phone, can be used to enable the user to copy an advertisement to his todo list or calendar, or be directly navigated to the place where he can buy the product. Second, a bonus coupon program could be used to trace back which advertisements led to which buys. The user could pick up a coded coupon at the advertisement, for example by taking a photo with his mobile phone, to get a discount at the shop. This coupon would encode the time and location of the advertisement. In addition to the feedback problem, we would have to deal with the cold start problem. For a new user, no information is available to estimate the mentioned probabilities. Therefore, we propose to use the m-estimate to start for each user with a global estimate of a mean user and gradually refine it as individual data is gathered. To estimate for example $P(f_j|ad, act_i)$, let act_{i,f_j} be the number of actions user i took with the feature f_j, act_i the total number of actions for user i, act_{f_j} the total number of actions with feature f_j for all users and act the total number of actions for all users. Let $P(f_j|ad, act) = \frac{act_{f_j}}{act}$. Then the m-estimate for $P(f_j|ad, act_i)$ is $\frac{act_{i,f_j} + mP(f_j|ad,act)}{act_i + m}$. M is also called equivalent sample size and can be interpreted as how often we assume to have observed user i to behave like the average user. A reasonable number for m would be 10, for example.

5 Conclusion and Future Work

We presented a simple approach to adapt advertisements on digital signage to the interests of the audience. We used auctions to sell the advertising space to the highest bidding advertisement and decision theory to determine for each advertisement how much to bid. Within this framework, we used methods from content based recommender systems to adapt the bid to the interests of the audience. Each advertisement has a set of keywords, and the history of all advertisements a user was interested in is kept. We proposed to use a naive Bayes classifier to estimate the probability that a user is interested in a certain advertisement given the keywords and the users history. We proposed interaction mechanisms to deal with the feedback problem and using the m-estimate to deal with the new user problem.
We plan to refine our method and evaluate other approaches from recommender systems for applicability to generate the bid as well. We plan to automate the feature identification, using expected information gain or the TF-IDF measure. In addition, we plan to evaluate our approach in a simulation as well as in a real-world experiment.

References

1. Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, June 2005.
2. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Working Paper, http://rwj.berkeley.edu/schwarz/, 2005.
3. Jörg Müller and Antonio Krüger. How much to bid in digital signage advertising auctions? In *Adjunct proceedings of Pervasive 2007*, 2007.
4. Kenton O'Hara, Mark Perry, and Elizabeth Churchill. *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*. Springer, 2003.
5. Terry Payne, Ester David, Nicholas R. Jennings, and Matthew Sharifi. Auction mechanisms for efficient advertisement selection on public displays. In *Proceedings of European Conference on Artificial Intelligence*, pages 285–289, 2006.
6. Michael J. Pazzani and Daniel Billsus. Learning and revising user profiles: The identification of interesting web sites. *Machine Learning*, 27(3):313–331, 1997.