How much to bid in Digital Signage Advertising Auctions?

Hans Jörg Müller* and Antonio Krüger*

Abstract. *In this paper we investigate bidding strategies for auctions that sell advertising space on digital signage. We introduce actionable advertisements, which are advertisements that entice you to take a specific action at a certain location in a specific time window. We introduce auctions for multiple advertising slots and present bidding strategies that take into account the number of faces seen in front of the display, the current time and location, as well as the user profiles of the users in front of the display. We use decision theory to calculate the expected utility of being shown for each advertisement.*

1. Introduction

Imagine yourself strolling around in a shopping mall of the future with your girlfriend in the evening. You already got the things you really need, and have a glimpse on one of the many digital signs as you pass by it. The sign presents you a special offer to get a coffee in the starbucks round the corner, and you are told that a celebrity gives an autogram hour at the other end of the mall, but you are both not interested. As you walk on, you are told that there is a life comedy starting in half an hour, but you don't feel like it. Then you see an advertisement that tells you that there is a piano concert in a piano store just 500 m away, and if you hurry up, you can get there just in time for the start of the concert and listen to a Sonata in F major by Mozart. Now that attracts your attention. As you both love Mozart, you rush to the piano store and have a great evening.
So how can such a digital signage be designed to deliver you the information you need to have a nice evening? Out of all the action possibilities available, it needs to present you those that have the highest probability to suit your needs, so you can scan these preselected items and choose the one you are really interested in. To this end, it should select the advertisments in dependence of the current context, for example the location, the time and your interests, when you opted in to be identified (e.g. via your Bluetooth enabled mobile phone).

2. Related Work

The concept of using auctions to select advertisements shown on digital signage was introduced in [5]. In contrast to this work, we focus only on actionable advertisements, like those mentioned in the introduction. Furthermore, we sell multiple advertising slots at the same time and use more context data to adapt the advertisement to the current context (camera, time, location, user profile). As auction mechanism, we use the generalized second price auction [1]. A general overview of work on situated public displays can be found in [4]. An early technical overview of our system can be found in [2].

*Institute for Geoinformatics, University of Münster, Germany, email: joerg.mueller@uni-muenster.de

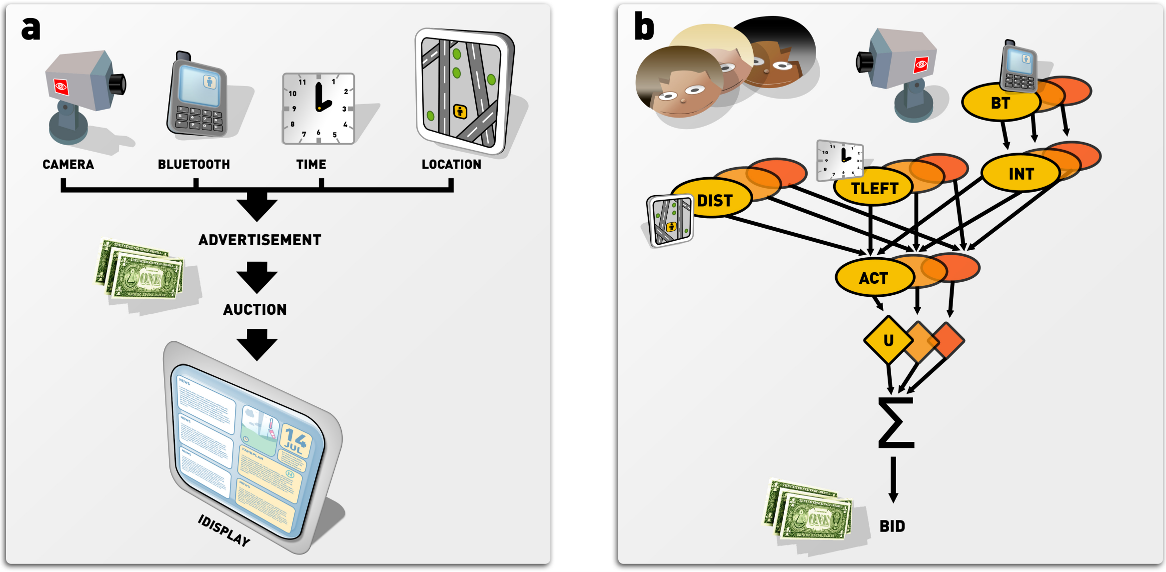

Figure 1. The auction process (a) and the bidding strategy (b) used in our approach. At the start of each advertising cycle, context information is gathered. The number of users looking at the display, if possible their identity, and the current time and location are provided to the advertisement. Advertisements take this context information and execute a Bayesian network to 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.

3. Auctioning for Actionable Advertisements

Actionable advertisements, like those mentioned in the introduction, entice you to take a clearly defined action within a certain window of space and time. The critical decision is whether you want to take the action or not. Thus, the task the user faces is the task of planning his time. From the total set of advertised action possibilities, the display selects some to present to the user. The user then has to determine what they are about and decide whether he wants to act. Once he decided, he could put the action to his to-do list or calendar, but in this paper we focus on the case where he wants to take the action immediately, possibly after he navigated to the location where he can take the action.
An important question is on what basis advertising space should be sold. Options would be per-impression, per-view and per-attendance. In per-impression, the advertiser pays for each second his advertisement is shown, regardless whether anyone has seen it. In per-view, the advertiser would pay for each person that has seen the advertisement. In per-attendance, the advertiser would pay for each person that has seen the advertisement and has then taken the action. Because those schemes are of increasing difficulty to implement technically, we chose pay per-impression to start with.
Similar to internet advertising on Google, we sell multiple advertising slots at the same time. These are displayed simultaneously next to each other on the advertising display, to increase the probability that the user is interested in one of the action possibilities. To start simple, we implemented a generalized second price auction similar to Google [1]. In difference to Google, in our case advertisers don't pay per-click, but per-impression. Because in the digital signage case we don't have keywords the user searches for, we provide context information to the advertisements to enable them to make better bids. The auction process is depicted in figure 1 a.

4. Smart Bidding

If advertising space is sold in an auction, the interesting question is how much an individual advertisement should bid to have itself shown. We employ decision theory to compute the expected utility of showing an advertisement or not $EU(AD)$. A Bayesian network is used to compute the corresponding probabilities. This approach enables us to smoothly integrate live sensor data and a priori estimations, obtained by marketing studies or mining the recorded data history. Thus, if in some occasion we cannot use a certain sensor, like the Bluetooth sensor, we can just use estimations of the a priori distribution of the corresponding variable. The bidding process is depicted in figure 1 b.
If we use an auction mechanism that has "truth-telling" as a dominant strategy [1], which we plan for the future, each advertisement should bid $EU(ad) - EU(\neg ad)$. These utilities depend on the profit the advertiser makes if the user takes the action or not $U(ACT)$ and the probability of the user taking the action $P(ACT|AD)$. We use upper case notation like $P(ACT)$ to denote probability distributions, while for binary variables lower case notation like $P(\neg act)$ denotes the individual probability that the variable takes this value. For simplicity we assume that the advertiser makes no profit if the user does not act $U(\neg act) = 0$ and that the client does not act if we do not advertise $P(act|\neg ad) = 0$. The set of users detected to observe the display is called $viewers$. The advertisement should then bid $EU(ad) = \sum_{i \in viewers} P(act_i|ad) * U(act_i)$. In order to take the context into account, we introduce three new variables $DIST$, $TLEFT$ and INT, that we assume to influence the probability. $DIST$ is the distance between the location of the action possibility and the display. $TLEFT$ is the time left until the user has to start walking to reach the location of the action possibility in time to take the action. If $TLEFT \leq 0$ the user can not reach that location in time anymore. Both values can be calculated from a graph model of the environment. An example of how such a model can be learned automatically is provided in [3]. INT is a binary value representing whether the user is interested in this kind of actions. We maintain a set of keywords for each user and each display. If at least one keyword matches, the user is assumed to be interested, otherwise not. Using these variables, we can refine our equation to:

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

We assume $U(act)$ is entered manually by the advertiser. If we assume that $DIST$, $TLEFT$ and INT are conditionally independent given ACT and ad, we can apply this and use Bayes' theorem to obtain:

$$P(act|DIST, TLEFT, INT, ad) = \frac{P(DIST|act, ad)P(TLEFT|act, ad)P(INT|act, ad)P(act|ad)}{P(DIST|ad)P(TLEFT|ad)P(INT|ad)}$$

We will now investigate each of the parameters in detail. $P(act|ad)$ can be approximated by the estimated number of participants divided by the estimated number of people who see the advertisement. $P(DIST|act, ad)$ is the probability that a participant has seen the advertisement at a certain distance from the action. $P(TLEFT|act, ad)$ is the probability that a participant has seen the advertisement a certain time before he needed to walk to the action possibility. $P(INT|act, ad)$ is the probability that a participant was considered interested in the action. All these parameters can be estimated manually or mined from history data. It would be very useful to have a way to find out where exactly the participants have seen the advertisement, for example by having the user pick up a voucher at the display and show it to the advertiser to get a discount. With such a kind of feedback, it would be easy to compute the mentioned probability distributions. $P(DIST|ad)$ and $P(TLEFT|ad)$ can be estimated by the percentage of advertisements shown at particular distances and with a particular time left. These percentages can be computed simply from a log of where and when the advertisement was shown. $P(INT|ad)$ is the percentage of people considered interested of those who see the advertisement. This can be approximated by the percentage of people who are interested, which is stored in the user profiles. Thus, advertisers have to provide for each advertisement at least how much it is worth to them that someone participates ($U(act)$), location, time, estimated number of participants and a set of keywords. With this data, the advertisement automatically bids the right amount in each auction, and the advertisement is shown to the right people in the right situations.

5. State of Implementation and Future Work

We have installed eight public displays at the University of Münster that run the system. We estimate the number of viewers either by the number of faces detected in front of the display or by the number of present BT devices, depending on whether a camera is available. Faces are only detected in frontal view, so each face found can be assumed to look towards the display. For those users who opted in to be identified, we identify them by the address of their BT enabled mobile phone. Because we provide a software with which people can interact with the display via their mobile phone, people are enticed to register their addresses to the display and provide user profiles associated with these addresses. Faculty can enter advertisements via an online interface similar to that used by Google AdWords. The whole system is implemented in Java, except for the face detection, which uses OpenCV and is implemented in C++.

As next steps we plan to improve the bidding strategy and the user model. We plan to take more context information into account and refine the methods for parameter estimation. In addition, we work on the interaction to entice more people to participate in our system and plan to evaluate our approach.

References

[1] 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.

[2] Hans J. Müller and Antonio Krüger. Towards situated public displays as multicast systems. In *UbiqUM 2006 Workshop on Ubiquitous User Modeling, The 17th European Conference on Artificial Intelligence*, 2006.

[3] Hans J. Müller, Oliver Paczkowski, and Antonio Krüger. Situated public displays that learn the structure of a building. In *Ubigis 2006 Second International Workshop on Ubiquitous GIS, GIScience 2006.*

[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.