Maximizing the Utility of Situated Public Displays

Jörg Müller[1], Antonio Krüger[1], and Tsvi Kuflik[2]

[1] Institute for Geoinformatics, University of Münster, Germany
[2] MIS department, University of Haifa, Israel

Abstract. Situated public displays are intended to convey important information to a large and heterogeneous population. Because of the heterogeneity of the population, they may risk providing a lot of irrelevant information. Many such important information items presented on public displays are actionables, items that are intended to trigger specific actions. The expected utility that such actionables have for a user depend on the value of the action for the user. A goal should be to provide for each user the actionables with highest utility. This can be achieved by adapting the information presentation to the users currently in front of the display. Adaptation can take place either by identifying individual users, by using statistics about the user groups usually in front of the display or by a combination of both. We present a formal framework based on decision theory that enables the integration of sensor data and statistics and allows to choose the optimal actionable to present based on this data.

1 Introduction

The falling costs of large displays and their potential usefulness increases the numbers of available public displays and they are starting to appear in many public places. Situated Public Displays (SPD) [5] are intended to convey useful information to large and heterogeneous populations, assuming that even though the characteristics of the users are partially unknown, the information provided on the SPD may be useful in a given context. Obvious examples are dynamic timetables at train/bus stations that present the planned timetable and any relevant updated information such as changes of platforms and delays. This information may be augmented by weather information at various destinations and even by the list of open coffee shops for passengers of delayed trains. Our previous research on traditional public displays such as paper-based pinboards and placards, has shown that most of the information presented to the public are actionables, which are intended to cause people to act. In the context of an university examples for actionables are talk announcements, open positions or special bargains for students. The associated actions can only be taken within a certain window of space and time, as for example talks take place in a certain room at a certain time and job postings are outdated after a while. In this work we are primarily concentrating on the presentation of actionables through

SPD. Even though the basic nature of SPD is to provide public information to large and highly heterogeneous populations, the basic assumption of complete anonymity may not be true. The user population in different places might be well known. For instance at universities there may be students, faculty and administrative staff with their particular interests and information needs. Hence actionables presented on public displays may be adapted to the interests of the current viewers and by that try to maximize their value for the users in question. The goal of this research is to establish a formal framework that will help estimating the expected utility of an actionable for users and groups alike. We are investigating the interesting question how to make use of different types of information about viewers and their interests, to finally decide which actionable to present on a SPD at a given moment in time.

2 Related and Prior Work

Research on various aspects of SPD has received considerable attention recently. SPD have been used to enhance the access for members of an organization to personal information anywhere within the organization, such as the BlueBoard system from IBM [7]. Additional work has been done to support spontaneous interaction between members of an organization, e.g. in the context of GroupCast [4], which aimed at improving interaction not directly related to the usual office work and supported social interaction by displaying mutual interests and hobbies on nearby large displays. Research on the CWall System [3] revealed the relevance to support groups of peers or Communities of Practice within organizations. Related to our work is also the Plasma Poster Network [2]. Here, the displays resemble real poster boards where anyone could post items to distribute information to people within the organization. The Lancaster ecampus project [9] is a campus wide installation of networked displays where several experiences have been made with displays at various locations in different contexts. Because of the more public nature of the installation, one important observation that has been made, is that the quality of content is very important and that deployment and maintenance costs should not be underestimated. A few longitudinal studies have looked at social and technical requirements of semi-public displays in organizations, such as door displays [1] or conference room reservation and notification tools [6].

In our department we have installed a system of five SPD presenting relevant information for students and faculty. The displays mainly show actionables and changes to actionables. The question that we will discuss in the following, is how to decide which actionable to present to which users, depending on the characteristics of the actionable and the interests of the users.

3 Deciding which Actionables to Advertise

Assuming the user in front of the SPD is identified (by a personal Bluetooth device for instance), the SPD content can be personalized by taking into account personal characteristics of the user (by having access to an individual user profile and schedule). As an example, the system may decide to remind an interesting talk given to a specific audience that seems to be of interest to the user even though he is not officially part of the target audience. Let us formalize the decision of selecting a single actionable to advertise.

We assume that there is a set of actionables $\mathcal{A} = \{a_1, \ldots, a_n\}$ that can be advertised and a set of users $\mathcal{U} = \{u_1, \ldots, u_p\}$ that can act upon these actionables. We have some evidence $\mathcal{E}$ that describes the situation that we can base our decision on, like sensor data, time, location etc. Then for each user u and each actionable a_k there is a utility $U_u(actupon(u, a_k))$ describing how useful it would be for that user to do the action, for example attend the talk.

An algorithm for finding the optimal actionable to advertise would now calculate the expected utility for each available actionable and choose the one that achieves the maximum expected utility. The expected utility for the user u $EU_u(advertise(a_i)|\mathcal{E})$ of advertising the actionable a_i given evidence $\mathcal{E}$ is what we expect to be the total utility for the user if we advertise this actionable. Following decision theory [8], we define

$$EU_u(advertise(a_i)|\mathcal{E}) = \\ \sum_{a_k \in \mathcal{A}} P(actupon(u, a_k)|advertise(a_i), \mathcal{E}) U_u(actupon(u, a_k))$$

Where, $P(actupon(u, a_k)|advertise(a_k), \mathcal{E})$ is the probability that user u takes the action a_k given that we advertise a_i and evidence $\mathcal{E}$. $U(actupon(u, a_k))$ again is the utility for the user u of taking the action a_k.

In reality, we may have more than one user in front of the display. Then we will have to maximize the expected utility $EU_{\mathcal{D}}$ for the whole group $\mathcal{D} \subseteq \mathcal{U}$ of users that are in front of the display. We assume that the probabilities of users taking an action are pairwise independent, and that the utility for the group is the sum of the utilities of the users. Thus, we can state:

$$EU_{\mathcal{D}}(advertise(a_i)|\mathcal{E}) = \\ \sum_{a_k \in \mathcal{A}} \sum_{u_l \in \mathcal{D}} P(actupon(u_l, a_k)|advertise(a_i), \mathcal{E}) U_{u_l}(actupon(u_l, a_k))$$

In a real-world setting, having all users identified may not be realistic. However, it could be possible to gather some statistics about which user groups usually pass the display at certain times. In our department for example, SPD are located at the entrances of and throughout the building. There are four different institutes each with three to seven groups of researchers. Each institute has students from five different years. In addition, there is administrative staff. Altogether there are more than forty different groups of users (with varying sizes) in that building. Information delivered over the public displays may be relevant to all (a change in the opening hours of the cafeteria) or parts of group members (a talk scheduled in one of the research groups). The different groups usually stay within certain regions of the building, and do so at different times. Students show up before first class starts, so no point in displaying schedule changes at

08:00 am when administrative staff arrives etc. Based on this idea, instead of dealing with individual users, we may take a stereotypic user modeling approach based on the characteristics of the groups (thus, we do not need to know the individual utilities $U_u(actupon(u, a_k))$). Let us suppose that there is a number of groups $\mathcal{G}$ such that each user is a member of one group. For each group $g \in \mathcal{G}$ we have an estimation $U_g(actupon(g, a_k))$ of how useful that action would be for members of the group. Then, we only need for each user u_l in front of the display and each group $g_m \in \mathcal{G}$, some estimation $P(u_l \in g_m|\mathcal{E})$ of the probability that the user belongs to this group, for example based on the current time, location of the display or sensor data. With this approach, the expected utility for the group would be

$$EU_{\mathcal{D}}(advertise(a_i)|\mathcal{E}) = \sum_{a_k \in \mathcal{A}} \sum_{u_l \in \mathcal{D}} \sum_{g_m \in \mathcal{G}} P(u_l \in g_m|\mathcal{E}) P(actupon(g_m, a_k)|advertise(a_i), \mathcal{E}) U_g(actupon(g_m, a_k))$$

Thus, we have presented a formula that calculates the expected utility of advertising an actionable given that we know a number of parameters.

4 Discussion

In the above we introduced the need for adaptation of public displays, the information sources that may be available for that and a formal definition that allows the selection of the best actionables to advertise, in order to optimize the users utilities. Two main issues become now the focus of our interest. The individual utilities $U_u(actupon(u, a_k))$ and the probability that a user will act upon an advertised actionable $P(actupon(u_l, a_k)|advertise(a_i), \mathcal{E})$. The utility of the action itself can be modeled as $benefit(u, a_k) - cost(u, a_k)$. The individual benefits are highly user dependent and should be estimated based on a user model. A user model may be represented in different ways: one way is a user model composed of weights assigned to concepts drawn from an organizational ontology (so every user has his/her own personal preference with respect to the common ontology). Another approach may be by a weighted vector of terms drawn from a domain vocabulary (or several domains). Other relevant aspects may be organizational role, education, marital status, age, preferences with respect to leisure activities and more. The "cost" side of the utility is less user depended and can be calculated e.g., considering time required, budget to be spent, and traveling distance. Assuming that we have the models for dealing with individual users, then group models can be calculated as an average of the individual models of the users that are members in these groups. The other way around, individual user models could be bootstrapped by group models if available. In addition to the individual utilities, we need an estimate of the probability of users acting upon the information. This may be impacted by the overall utility value, the time left until the deadline, the possible alternatives, the need to change existing plans and so on. A strength of our approach is that some of the parameters can be estimated online from sensor data, while others can be obtained from statistics. So it adapts easily to different situations where different sensors and amounts of a priori knowledge are available, and it is does not matter that some users are identified and while others are not. A weakness of our formalism is that each user can be member of only one group, but this can easily be circumvented by modeling intersections of groups as additional groups. A different aspect to be dealt with is the optimization of the limited space and time that is available for all possible information items on a SPD. Of course we only take into account the utility for users, so we do not look at the particular needs of other stakeholders like information providers or display owners, who themselves might have their particular interests and want the information to be presented only to certain groups of individuals.

Acknowledgements We would like to thank Oliver Paczkowski and our team for all the work on the display environment and Martin Mühlenbrock for hints on advertising.

References

1. K. Cheverst, D. Fitton, and A. Dix, 'Exploring the evolution of office door displays', in *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, eds., K. O'Hara, M. Perry, E. Churchill, and D. Russell. Kluwer International, (2003).
2. E. Churchill, L. Nelson, L. Denoue, and A. Girgensohn, 'The plasma poster network: Posting multimedia content in public places', in *Human-Computer Interaction INTERACT 03*. IOS Press, (2003).
3. Antonietta Grasso, Martin Mühlenbrock, Frederic Roulland, and Dave Snowdon, 'Supporting communities of practice with large screen displays', in *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, eds., K. O'Hara, M. Perry, E. Churchill, and D. Russell, Kluwer International, (2003).
4. Joseph F. Mccarthy, Tony J. Costa, and Edy S. Liongosari, 'Unicast, outcast & groupcast: Three steps toward ubiquitous, peripheral displays', in *Proc. Ubicomp 2001, LNCS 2201*, pp. 332+, (2001).
5. Jörg 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).
6. Kenton O'Hara, Mark Perry, and Simon Lewis, 'Situated web signs and the ordering of social action', in *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, eds., K. O'Hara, M. Perry, E. Churchill, and D. Russell, Kluwer International, (2003).
7. Daniel M. Russel and Alison Sue, 'Large interactive public displays: Use patterns, support patterns, community patterns', in *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, eds., K. O'Hara, M. Perry, E. Churchill, and D. Russell, Kluwer International, (2003).
8. Stuart J. Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach (2nd Edition)*, Prentice Hall, 2002.
9. Oliver Storz, Adrian Friday, Nigel Davies, Joe Finney, Corina Sas, and Jennifer Sheridan, 'Public ubiquitous computing systems: Lessons from the e-campus display deployments', *IEEE Pervasive Computing*, **5**(3), 40–47, (2006).