Beyond Web 2.0: Challenges in Personalizing for Networked Public Display Environments

Nemanja Memarovic, Marc Langheinrich

Faculty of Informatics, University of Lugano (USI), Switzerland
{nemanja.memarovic,marc.langheinrich}@usi.ch

Abstract. The falling prices of LCD technologies have led to a proliferation of large public displays. Today, most of these displays show simple slide shows or loop through a set of movies in isolation, yet efforts are underway to network these displays together on a global scale. Networked public display would go beyond simple ad-pushing and open up towards truly personalized services, such as support for navigation tasks, messaging, emergency services, and of course context-aware advertising. Moving from personalized Web content to personalized public display content, however, leads to novel challenges that limit the applicability of existing personalization solutions. This article enumerates the key research challenges that stem from bringing personalization to public displays, namely in the areas of user identification; profile location; profile content; content tailoring; model refinement; and applications that require personalisation.

1 Introduction

Public displays are becoming ubiquitous in our environment: we can find them on buildings, in shopping malls, universities, museums, buses, and in train stations. While many of those displays are still singular installations that simply run some slideshow off a locally connected computer, it is easy to envision that in the future, all of these displays would be connected to the Internet, thus constituting a novel global communication medium. An interconnected global network of public displays would allow for a range of novel, personalized and context-sensitive applications, such as:

- **Ambient personalized navigation:** A public display in the lobby of a concert hall could indicate bus and tram departure times, highlighting and indicating those that run in the direction of the homes of the people coming out of the concert. An in-store display could list directions to individual store sections based on the shopping interests of the people nearby.
- **Personal public messaging:** Instead of writing on virtual "walls" inside a Facebook account, a public display at an intersection would offer truly wall-sized messages targeted at individuals passing it, offering, e.g., encouragement for today's exam ("Go John, you can do it!") or public displays of affection ("I love you Linda!").

– **Context-aware advertising:** A long standing dream of bricks-and-mortar retailers, networked public displays could configure in-store signage to optimally suite the current shoppers' interests. This can be especially interesting for small neighborhood shops that can help each other out by unobtrusively steering customers with matching needs to a shop down the road.

In each of the examples above, individual or even group-based personalization is key to offering highly useful information and services on the hithereto largely ignored displays [9]. Although personalization is used extensively on the web, where content is regularly customized to fit individual users or even groups of users, integrating personalization technology into networked public displays is not as simple as installing a web browser on the public display.

Even the first step in personalization – user identification – is difficult to carry over from the web into networked public displays. Web systems can draw on explicit authentication mechanisms such as usernames and passwords, cookies, or at least IP addresses. How could this be supported in a public display network setting? Manually entering credentials seems awkward at best; using the user's mobile phone to store and replay some cookie-like identifyer looks prone to security and privacy issues. And what if the user is not carrying a phone, or if the user's phone cannot be identified due to compatibility issues? While one could use a camera for a visual identification instead, face-recognition technology is still far from being able to reliably identify arbitrary individuals in non-lab situations, not to mention the lack of public acceptance for such uses.

With similar challenges along the entire user-display interaction cycle, we will need to come up with novel approaches to personalization when we want to move interaction from personal desktops into public spaces. This article enumerates and explores these new challenges.

2 Background

The field of networked public displays has recently emerged from a number of related fields, most notably the work on situated public displays [18] and the young field of pervasive advertising [15, 17].

Individual personalization in public displays so far has mainly relied on Bluetooth scanning to detect the presence of registered users in the vicinity of a display. Strohbach et al. [21] are building a context-aware public display service platform in their "Pervasive Display Networks (PDN)" initiative. Their initial prototype system demonstrates how various system services work together to acquire and distribute generic context information, e.g., to display individual product recommendations on the public displays or to show targeted product discounts (in the form of a barcode coupon displayed on-screen). Registered users are identified by the Bluetooth MAC addresses of their mobile phone. Profile acquisition, representation, and (non-trivial) production,[1] however, have so far been outside the scope of their work.

[1] Kobsa et al. [12] characterise a (web) personalization process to be comprised of three parts: (1) acquisition, i.e., gathering information about the user and building a

José et al. [10, 11] and Davies et al. [6] have experimented with using Bluetooth *device names* to personalize public display, alleviating the need for user registration. In several fielded experiments, users could put simple commands into their mobile phone's Bluetooth name (e.g., "*tag:punk,pop*" to indicate music preferences, or "*flickr:oranges*" to trigger a flickr search for pictures of oranges). Besides the privacy-friendly nature of not requiring a registration, this method also is free and supports hands-free operations. Initial trials in a bar in Portugal and on a University Campus in the UK have received very positive feedback. However, this approach clearly does not permit long-standing user modeling nor more complex services that go beyond simple search keywords.

Müller and Krüger [16] use Bluetooth scanning to model user movement among a network of public displays, thus learning the actual topology of the displays in order to coordinate shows among multiple displays that "follow" the user, thus increasing exposure. This example demonstrates how, instead of adapting to the user's situation, networked public displays can also adapt to the *situation* they are in, e.g., the number of people in front of it, or where they came from and where they are going. Cardoso and José [3] offer a generic context adaption framework for such situated public displays. Bluetooth scanning is just one attribute of the "*digital footprint*" of a display, allowing it to react to a range of contextual parameters, such as the number of "eyeballs" that are currently watching it [14]. This work is highly relevant, as it shows that personalization in public displays can and must take the "big picture" of the actual deployment into account. Personalized networked public displays need to combine the context of a place and its changing audience with the individual information needs of each user.

Current research efforts in personalization technology are still mostly targeting web environments (see, e.g., Gauch et al. [8] for an overview), though context-aware systems have recently been gaining in popularity (see, e.g., Baldauf et al. [2]). Assad et al. [1] have proposed a distributed context model called *PersonisAD* that allows applications to discover and use distributed models of places and users in order to tailor their services. While the generic nature of PersonisAD could be easily adapted to work with networked public displays, Assad et al. do not address access control and distributed ownership issues, but assume a single provider with a full registered user base.

Privacy is a major issue in personalization research, and much work has proposed mechanisms to provide privacy-enhanced personalization, both for traditional web applications [4, 22] and recently also for mobile environments: Kuflik and Poteriaykina [13] suggest making the user's mobile phone act as a mediator between the user's protected user model (stored somewhere online) and the local service that is looking for model data in order to personalize its offering. Darke et al. [5] augment the PersonisAD architecture with privacy rules that

model; (2) representation, i.e., expressing the user model and performing inferences; and (3) production, i.e., the actual adaption of content and services based on the user model.

can govern the sharing of user models across services. It remains to be seen how well these approaches actually work in a networked public displays setting.

3 Challenges

As stated above, personalizing networked public displays will require more than just transferring the techniques that are already in use on the web. We see six key research challenges, which are summarized in Figure 1: the process of identifying of the user; the proper location of the user model; the actual content of the model; the effective tailoring of the display; the efficient refinement of the model; and finding suitable applications that would use the model. We will discuss these challenges in turn below.

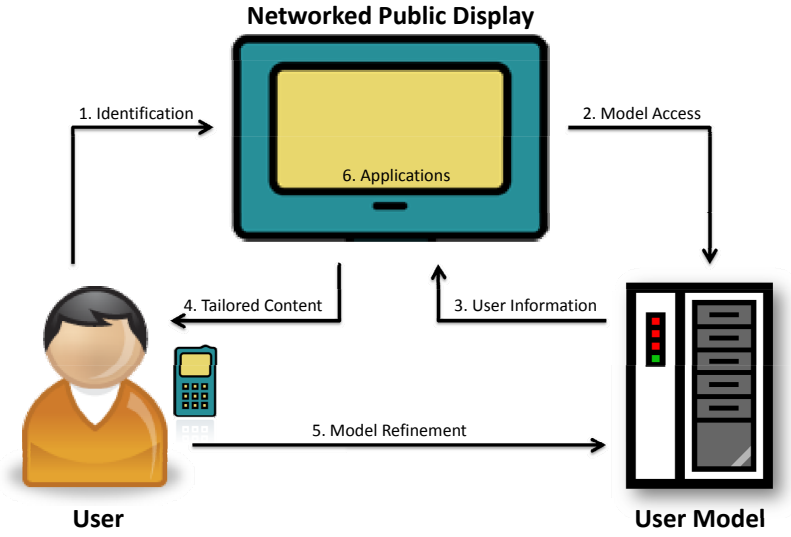

Fig. 1. Challenges for personalising networked public displays: 1) User identification 2) User profile/model location 3) Profile content 4) Content tailoring 5) Model refinement 6) Selection of suitable applications and their personalisation

3.1 User Identification

The first challenge that has to be answered is how can we identify the users. What *method* should we use for identification: face detection, eye gaze gestures, biometrics, some sort of identification via the user's mobile phone, or something else? There might not be a single right answer to this question. Also, what is the *extent* of user identification in the networked public display environment?

When a user is identified in the system, is he/she identified for that particular display only or is there a need for some sort of single sign-on? And how long should a user be identified within the system? One of the novelties that has to be considered when thinking about the user identification with the networked public displays is the *surrounding* of the display's environment [3]. The user's distance to the display could be one of the properties of display surroundings that would influence the identification method. For example, face detection would work well with the displays that are near the users, e.g., a display located in a public library, but it would not work that well for the displays that are located far away from the users, e.g., a billboard-size display at a public square. Another property of the display's surrounding could be traffic, i.e., the number of people that are near the display. Is there a need to adjust identification techniques depending on the number of people that are around the display?

3.2 User Model Location

Once a user is identified at a display, his or her profile has to be retrieved. But retrieved from where? User profiles could be *distributed* across the entire network and then merged at each display as needed [1]. The profile could also be in a *cloud*, i.e., distributed but with a single point of entry, resembling a centralized system from the administrative point of view. The profile could also be *roaming*, i.e., pulled on-demand from some user-designated home location (as in [13]), or *portable*, i.e., downloaded directly off a mobile device that the user carries.

3.3 Profile Content

After the profile is retrieved it has to be read so that the content of the display can be personalised. But what information is stored inside the profile? What information is relevant for personalisation with the networked public display? We can start from some standard user demographics, e.g., age, gender, race, but how much of the user's history is relevant? And do we also have to consider additional user context such as the time of day, the display's location, traffic, and the user's social setting (is he alone, with friends, family)? Is there a need for modeling user profiles across domains [19]? Or should each application maintain their own, separate model? One option might also be the re-use of existing web profiles, such as user-maintained profiles on LinkedIn, MySpace, or Facebook [20].

3.4 Content Tailoring

With personalized content being displayed in a public or semi-public setting, the actual presentation of tailored content has an added twist in a networked public display setting. Should the personalisation be *subtle* or *noticeable* for the user? Subtle personalization lowers the risk of embarassing users with custom-tailored content. Subtle display changes also work better when the display needs to be adapted to multiple users. However, too subtle a change and users might not

notice (nor appreciate) the personalization. More noticable tailoring could offer more tangible benefits for each user, with a lower likelihood that useful information is overlooked. Maybe the question really is not about *if* personalization should be subtle or noticeable, but *when* it should be subtle and when it should be noticeable. For example, in an advertising context, subtle displays might work better than explicitly addressing individuals with offers, unless a user explicitly demonstrates that she is looking for a particular offer (then a highly personalized ad might be much more effective). On the other hand, a public display network in an airport might opt for explicitly targeting late passengers for boarding calls, in order to more effectively direct them to the right gate.

3.5 Model Refinement

Today's web based personalization system can keep detailed tabs on the user's interaction with the system with the help of page views, mouse clicks, and even purchase records. How is this going to be done with networked public displays? *How are we going to learn about users?* Modern smartphones could be used to *record* user's actions – be it with the display itself or maybe even beyond the immediate display interaction. Alternatively, web interactions could be re-used through the integration of web based profiles (see above). However, with such an inclusive view to model refinement, users might need to explicitly define boundaries to these profiles – either in a static fashion (i.e., "only incorporate Amazon and Facebook interactions into my profile") or in a dynamic fashion (i.e., "only integrate my interactions with a public display if more than 10 people are around me – do not record my more intimate interactions"). Another important question is *how are we going to make sense of a user's actions?* For example, a user passes an advertisement showing a product that fits into his "products of interest" list from his profile, and it is on sale in a nearby store. What happens if the user does not end up following-up on that sale? Was it because the ad was ignored or not seen? What is the alternative to the web's "click-through" feeback? Eyeballs [14], i.e., the number of people that looked at the advertisement? We have to enlarge our model of user interactions once we move from desktop interactions to real-world interactions, in order to better interpret potential positive and negative feedback to public display personalization.

3.6 Application Selection

Personalisation on the web comes in the form of customized on-line stores, tailored information portals, or individual communication sites such as social networking platforms. What are the "killer-apps" for personalized networked public displays? Information and infotainment certainly should work well, and personalized advertisements are of course the prime economic driver for industrial interest in this field. How can we go beyond news stories and ad tailoring? Can we transform today's simple one-way displays into a new bi-directional communication medium, based on novel personalization techniques that offer effective, efficient, and enjoyable interactions – with service providers, advertisers, and other users?

4 Conclusion

Today, most public displays are stand-alone installations, yet soon we might find large numbers of them networked together. Personalizing content for such networked public displays is an intriguing value proposition that could radically change public spaces: from todays environments in which information is pushed to passers-by in the form of adverts, to spaces that can utilize public displays and ambient intelligence to reflect the hopes, aspirations and interests of its occupants using content and applications created anywhere in a global network. While a large body of research exists in the field of web personalization, transferring these techniques is limited by the novel challenges that networked public display environments pose. In this paper we described six key challenges that we believe need to be addressed in order to bring effective personalization into this novel domain: (1) *user identification*, i.e., how is the user going to be identified in the system; (2) *user profile location*, i.e., where is the user profile going to be stored; (3) *profile content*, i.e., what information should be stored in user's profiles; (4) *content tailoring*, i.e., how will the users know that the content of the display is personalised just for them; (5) *model refinement*, i.e., how is the system going to learn about the users; and (6) *applications that require personalisation*, i.e., what applications for the networked public display environments require personalisation.

Acknowledgements

The challenges of networked public displays have been the focus of the Dagstuhl Seminar 10011 on "Pervasive Public Displays" [7]. We thank all participants for the fruitful discussions surrounding pervasive public display personalization.

References

1. M. Assad, D. J. Carmichael, J. Kay, and B. Kummerfeld. PersonisAD: Distributed, active, scrutable model framework for context-aware services. In A. LaMarca, M. Langheinrich, and K. N. Truong, editors, *Pervasive Computing, 5th International Conference, PERVASIVE 2007, Toronto, Canada, May 13-16, 2007, Proceedings*, volume 4480 of *LNCS*, pages 55–72. Springer, 2007.
2. M. Baldauf, S. Dustdar, and F. Rosenberg. A survey on context-aware systems. *Int. J. Ad Hoc Ubiquitous Comput.*, 2(4):263–277, 2007.
3. J. C. S. Cardoso and R. José. A framework for context-aware adaptation in public displays. In *On the Move to Meaningful Internet Systems: OTM 2009 Workshops*, volume 5872 of *LNCS*, pages 118–127, 2009.
4. I. Cingil, A. Dogac, and A. Azgin. A broader approach to personalization. *Commun. ACM*, 43(8):136–141, 2000.
5. G. Darke, J. Kay, and B. Kummerfeld. Mobile personalisation: new challenges for privacy. In *Proceedings of the 7th International Workshop on Ubiquitous User Modelling (UbiqUM2009)*, pages 41–44, 2009.

6. N. Davies, A. Friday, P. Newman, S. Rutlidge, and O. Storz. Using bluetooth device names to support interaction in smart environments. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 151–164, New York, NY, USA, 2009. ACM.
7. N. Davies, A. Krüger, M. Langheinrich, A. Schmidt, and M. Strohbach. Pervasive public displays – Dagstuhl seminar 10011. http://tinyurl.com/yglgtc4, January 2010.
8. S. Gauch, M. Speretta, A. Chandramouli, and A. Micarelli. User profiles for personalized information access. In *The Adaptive Web*, volume 4321 of *LNCS*, chapter 2, pages 54–89. Springer, Berlin, Heidelberg, 2007.
9. E. M. Huang, A. Koster, and J. Borchers. Overcoming assumptions and uncovering practices: When does the public really look at public displays? In J. Indulska, D. J. Patterson, T. Rodden, and M. Ott, editors, *Pervasive Computing, 6th International Conference, Pervasive 2008, Sydney, Australia, May 19-22, 2008, Proceedings*, volume 5013 of *LNCS*, pages 228–243. Springer, 2008.
10. R. José and F. Bernardo. Extended Bluetooth naming for empowered presence and situated interaction with public displays. In J. M. Corchado, D. I. Tapia, and J. Bravo, editors, *Proceedings of the 3rd Symposium of Ubiquitous Computing and Ambient Intelligence 2008*, pages 57–65, 2009.
11. R. José, N. Otero, S. Izadi, and R. Harper. Instant places: Using bluetooth for situated interaction in public displays. *IEEE Perv. Comp.*, 7(4):52–57, 2008.
12. A. Kobsa, J. Koenemann, and W. Pohl. Personalised hypermedia presentation techniques for improving online customer relationships. *The Knowledge Engineering Review*, 16(02):111–155, 2001.
13. T. Kuflik and K. Poteriaykina. Integrating user modelling server with user modelling mediator on a personal device. In *Proceedings of the 7th International Workshop on Ubiquitous User Modelling (UbiqUM2009)*, pages 11–15, 2009.
14. J. Müller, J. Exeler, M. Buzeck, and A. Krüger. Reflectivesigns: Digital signs that adapt to audience attention. In *Pervasive '09: Proceedings of the 7th International Conference on Pervasive Computing*, pages 17–24, Berlin, Heidelberg, 2009. Springer.
15. J. Müller, P. Holleis, A. Schmidt, and M. May, editors. *Proceedings of the 2nd Workshop on Pervasive Advertising*, volume 4480 of *LNCS*, Lübeck, Germany, October 2009. Springer.
16. J. Müller and A. Krüger. Learning topologies of situated public displays by observing implicit user interactions. In *Proceedings of HCII Human Computer Interaction International 2007, Beijing*, 2007.
17. J. Müller, A. Schmidt, A. Quigley, and B. Begole, editors. *Proceedings of the 1st Workshop on Pervasive Advertising*, volume 4480 of *LNCS*, Nara, Japan, May 2009. Springer.
18. K. O'Hara, M. Perry, and E. Churchill. *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*. Kluwer Academic Publishers, Norwell, MA, USA, 2004.
19. J. Schuurmans and E. Zijlstra. Towards a continuous personalization experience. In *Proceedings of the conference on Dutch directions in HCI*, page 19, New York, NY, USA, 2004. ACM.
20. R. Shannon, M. Stabeler, A. Quigley, and P. Nixon. Profiling and targeting opportunities in pervasive advertising. In J. Müller, A. Schmidt, A. Quigley, and B. Begole, editors, *Proceedings of Workshop on Pervasive Advertising*, pages 42–45, Nara, Japan, May 2009.

21. M. Strohbach, E. Kovacs, and M. Miquel. Pervasive display networks - real-world internet services for public displays. In *Adjunct Proceedings of the 4th European Conference on Smart Sensing and Context (EuroSSC)*, pages 39–42, 2009.
22. Y. Wang and A. Kobsa. Technical solutions for privacy-enhanced personalization. In C. Mourlas and P. Germanakos, editors, *Intelligent User Interfaces: Adaptation and Personalization Systems and Technologies*, pages 353–376. Information Science Reference, 2009.