

Timely and keyword-based dynamic content selection for public displays

Fernando Reinaldo Ribeiro

Informatics Department
Polytechnic Institute of Castelo Branco
6000-767 Castelo Branco, Portugal
fribeiro@est.ipcb.pt

Rui José

Information Systems Department
University of Minho
4800 Guimarães, Portugal
rui@dsi.uminho.pt

Abstract— In public display systems determine what to present and when is a central feature. Although several adaptive scheduling alternatives have been explored, which introduce sensibility of the display to some type of external variable, they are still very dependent on the user in their behavior, content specific in their nature and very rigid in their adaptation to their social environment, not providing visitors of the place with appropriate, rich and personalized information according to their interests and expectations. There is a need for solutions that successfully integrate the wealth of dynamic web sources as providers for situated and updated content with social and contextual environment around the display so as to present the most appropriate content at every moment, and thus improving the utility of the system. In this paper, we present a recommender system for public situated displays that is able to autonomously select relevant content from Internet sources using a keyword-based place model as input. Based on external relevance criteria the system finds and pre-selects only those sources that are more relevant, and an adaptive scheduling algorithm continuously select content that are relevant, timely, in accordance with the place model, sensitive to immediate indications of interest and balanced to serve the broad range of interests of the target population. To evaluate this system we have carried out two partial experiments. The results showed that keyword-based shared place models jointly with content specific relevance models are a simple and valid approach to user-generated content for public displays.

Keywords- *Public Displays, Situated Displays, Intelligent Environments, Information Integration, Social Information Systems, Web Information Filtering and Retrieval, Ubicomp.*

I. INTRODUCTION

Even though large public displays have always been part of the ubiquitous computing vision (the “boards”[1]), their use has attracted considerably more interest in recent years, given the increasingly pervasive presence of plasma screens, projectors and smaller TFT. These advances in screen technology have resulted in a much wider range of affordable displays becoming available, providing new opportunities for supporting collocated working, learning and socializing. However, despite this widespread use, studies have shown that most public displays are not very valued by their potential users [2]. Many of them are non-interactive and are typically used as mere distribution points for pre-packaged and centrally created content, resulting in institutional content that is perceived by users as too static and largely unrelated with the particular place where the display is installed.

These displays as public, shared and pro-active devices that are embodied into their environment should reflect the information and the interactions associated with that environment and the people on it, enabling new and more engaging user experiences sensing their environment, giving users a more active role in the system behavior, and providing people with brief encounters with information that is relevant for their specific situation. This could improve local awareness, enrich our perception of reality and ultimately make situated displays an integral part of the social settings in which they are integrated. To achieve this, public and situated displays should go beyond the presentation of pre-defined content through non adaptive scheduling processes or simply reacting to explicit users’ interactions. They should dynamically find and select the most relevant content according to the current context of presentation. In this scenario, the wealth of content and information sources on the web provides an obvious, undeniable and tempting solution to the problem of content management in public displays. Content from web content providers, like news, headlines, blogs, events, and many others, represent an important resource and can play a central role, providing regularly updated data available for use by the display. They can be leveraged for the generation of content for digital situated displays brings up displays as highly personalized devices to places where they are situated, presenting the most appropriate content according to its social environment. To really take advantage of the potential of web information, the display system should be able to dynamically discover and select from the Internet information sources those sources that are more relevant for each particular display and then, they should proactively select and present the most appropriate content at the right time according to its relevance to the current context of users and their environment. This would make each display system unique and closely related with the specific place where it is installed, providing the ground for highly situated displays that reflect the expectations, interests and practices associated with each particular place.

This work is particularly focused on situated displays designed for shared and communal use in public and semi-public settings. People in these places are mainly occasional visitors with sporadic or even single contact with the display and most people will not know each other. The integration of their individual preferences, many times characterized by distinct interests and expectations, into the same social context make this type of setting very dynamic and heterogeneous. To be adaptive to these changes the display

should have the ability to adapt to this continuous flow of social settings and the ability to contribute to the emergence of a shared concept of place that matches the broader social expectations and practices of the community as a whole. These abilities suggest approaches that build strongly on user participation. In particular, the system should enable high levels of appropriation by multiple communities of practice, possibly by focusing on basic forms of appropriation in which the creation of meaning is largely left to people.

The main objective of this work is to develop and evaluate a dynamic sources recommender system for public interactive displays that dynamically finds and selects the most appropriate content according to the social environment around the display. This broad research goal embraces the following set of research objectives:

a) Implement and validate a place model that represents the shared contributions of the place owner interests and place visitors interests and acts as guidance for the display behavior.

b) Implement and evaluate methods for obtaining relevant sources based on keywords from the place model, using dynamic web sources as content providers and relevance criteria for source pre-selection.

c) Implement and evaluate an adaptive scheduling algorithm that is continuously selecting content that is relevant, timely, in accordance with the place model, sensitive to immediate indications of interest and balanced to serve the broad range of interests of the target population.

The remainder of this paper is structured as follows: in Section 2 we discuss the related work. In Section 3 we present our approach to the problem of recommending relevant and timely dynamic content for public displays. Section 4 describes the evaluation of two partial prototypes in real scenarios and we also discuss obtained results. Finally, in section 5 we present general conclusions of this work and proposals for future work.

II. RELATED WORK

To provide a background, we will first survey existing work in two main areas: adaptive scheduling in public displays and recommender systems in general.

In public display systems determine what to present and when is a central feature. In most commercial systems, this is a fixed schedule based on a cycle of pre-defined content. Several adaptive scheduling alternatives have been explored that introduce sensibility of the display to some type of external variable. Proactive displays [3, 4] select content that is scheduled on-the-fly according to the interests of users within the direct vicinity of the display. Several types of interactive displays targeted at direct collaboration, such as the Dynamo [5] or BlueBoard [6], give users direct control of the display and thus content selection are typically resolved by users rather than by the system. BlueScreen [7] bases its scheduling decisions on history observation of exposed devices which were collected during the advertising cycles. Muller [8] proposes a system that uses a naïve Bayes classifier and information about history, current time,

location, advert keywords and feedback to determine the best advert to display. Even though some of these systems are able to support unassisted and adaptive scheduling they employ customized scheduling algorithms based on users interests or user profiles. In this work we take a generic place specification and we explore how it can be used as the basis for autonomous content selection. We assume that we have no a priori knowledge about users' profiles but a simple high-level specification of place. Furthermore, our approach is not specific to any particular type of content. Even though in this paper we put more focus on information feeds, the same specification could be used for selecting videos, photos, advertisements, or any other type of content.

Regarding previous work in recommender systems and information retrieval, we tried as much as possible to leverage on existing algorithms and tools. Although single content-based techniques have been used in news recommender system (e.g. Bogers et al [9]) it is our belief that this approach is not suitable for our problem. This is corroborated by Das, A.S., et al [10] that refers that content based approach is not suitable to domains such as news. Collaborative filtering approaches also present some drawbacks that are commonly associated to this technique that limit the applicability of this approach to our problem, namely: the system cannot recommend items that are different from anything the user has seen before and in many cases this is what happens in our work. In fact, the lack of users' feedback on past items, associated to the high number of content and the diversity of users on public displays environments, makes common approaches to recommender systems inappropriate for this scenario.

III. RECOMMENDING RELEVANT AND TIMELY CONTENT FOR PUBLIC DISPLAYS

To make the content selection adaptive we may follow several possibilities. One is to try to embed into the behavior definition, or directly into the scheduling algorithm itself, empirical knowledge that represents how we expect the system to be affected by changes in its context. For example, we may want to say that a particular piece of content should be shown only when there are Bluetooth devices nearby or that another item should only be shown when there is no one in the immediate physical proximity of the display. As these types of rules refer to contextual variables they can in fact enable dynamic scheduling in which the behavior of the system is continuously adapting to the varying context of its operating environment. However, such rules are difficult to define and represent, mainly because smartness is a rather vague concept that for most cases does not map directly into any specific type of reaction. For this reason, this approach is only efficient when the goal is to create a very direct association between a particular state and a very specific behavior, normally something in the form of a trigger.

Other possibility is to try to create new domain knowledge by training the system according to some general notion of smart behavior. This would typically include a training stage, in which a set of training cases of the expected behavior would be generated, a process of supervised learning, in which new rules would be inferred, and then a

dissemination stage, in which those rules would become domain knowledge ready to be embedded into scheduling processes. It is however unclear how much of that knowledge is generic enough to be applied to multiple situated displays, as the respective context and media items may be completely different.

Thus, in our approach we decide to put less focus on sensing, modeling and inferring context and more focus on people and their interactions and dynamic data about information items, as the main drivers for the adaptability of the system. Empowering place visitors and place owner to influence the content selection gives to the situated display a more social behavior reflecting the preferences of users that share the place and selecting from web sources content that would be highly relevant for that place.

We will now describe in more detail the operation of this recommender system, which is organized, as shown in Fig. 1, in 3 sub-systems: Place, Selector and Scheduler.

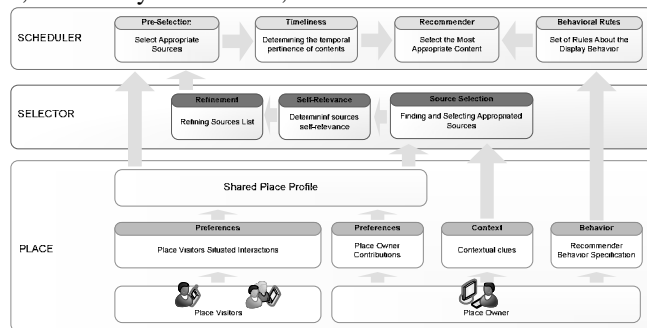


Figure 1. System architecture.

A. Place

The situated nature of the display means that content presented by the display is related to the physical space where the display is situated and related to people and activities that take place in that environment. Thus, the display should be aware of the place where it is installed, its characteristics and the interests and expectations of people that frequent the space. Recommendations should be performed considering the social environment around the display, which integrates the preferences of both place owner and place visitors into the same representation and thus, benefiting from synergies between them, serving the community as a whole instead of each individual interest at a time.

These environments are characterized with a high variability of place visitors in number and frequency, each one with specific interests and expectations, and thus, they are very dynamic and heterogeneous. Although individual contributions per se do not represent the place nature, the combination of various contributions along time, allows a characterization of the place social environment as a seamless sequence of many social settings. Thus, we choose to model this knowledge into a shared place model through a weighted list of keywords, a tag cloud, corresponding to topics of general interest for the environment where the display is placed, avoiding the usage of complex configurations and laborious models for describing the

preferences of the stakeholders, for example by asking for keyword pairs or by proposing ontologies, what makes this a burden and specialized task.

This allow place visitors and place owner to realize their contributions through simple and non intrusive keyword-based interactions and their integration into a shared place preferences model, representing synergies between them and also their temporal relation. This shared place model is central for the display behavior. It includes the preferences of all stakeholders in the same environment model and it is the base for source selection, i.e. the display uses this knowledge to find related sources and then to select the most appropriate content at each moment.

To support place visitors' interactions we follow the approach of empower place visitors by giving them a way to use their Bluetooth device name as a mean for self-expression and situatedness (as described in [11]). This approach go beyond traditional approaches that focus on Bluetooth scanning as a surveillance mechanisms were people are merely sensed and do not play any active role in the meaning of their exposure. It gives to the place visitors simple and unobstructive mechanisms to express their interests through simple keywords and thus, to contribute to the social context of the place

Place owner specifies his interests through simple keywords and configures the display behavior through simple rules. He also may contextualize the display environment giving more sense to situated interactions. This information about the context may considerably improve the selection process by framing place visitors' keywords in a broader place context defined by physical location, organizational context, type of setting, or even the overall set of user-generated keywords. Even simple contextual clues may be enough in disambiguating ambiguous keywords and in providing an interpretation aligned with the nature of place.

As user-generated keywords these contextual tags must be chosen appropriately in a way to maintain the situated behavior of the display. Using for example the scenario of the University of Minho located in the city of Braga, contextual information may be as example {braga, minho, education, university, ...}. This information is very important helping to give sense and situated connotation to the keywords obtained through users' interactions. If a user suggests the keyword "images" the meaning of the keyword is very diffuse and the display may present images about any topic. However, when contextualized with the previous contextual information the display may select images related to the location of the university or related to the university itself.

The usage of simple keywords for preferences description presents some advantages: they are simple to describe; they need simple interaction mechanisms, non intrusive and non hard work to perform; they are widely used, which makes it easier to extract them for all types of information objects and; all contributions from place visitors are simple to integrate in a shared place profile.

B. Selector

The main purpose of the selector sub-system is to autonomously find and select relevant content sources from the Internet that should be appropriated for the current social environment around the display.

The enormous quantity of data available from Web sources makes it an important resource for providing dynamic and updated content. Their usage as content providers for situated displays, allows surpassing two important problems. First, eliminate the need of finding and pre-defining a set of sources with relevant and situated content. Second, overpass the limitations of the usage of a reduced set of categories of content imposed by the usage of pre-defined list of sources.

Dynamic sources are very heterogeneous and dynamic, mainly due its frequently updated and diverse content, what makes difficult to know how relevant is the content that each source has to offer. More than to know information about source availability, source stability or contextual quality of the contents, it is important to know how relevant is the content that each source has to offer at each moment. Thus, to benefit from the high potential of these sources two aspects are central: finding those sources that are considered relevant for the current context and evaluate the temporal pertinence of its content. Only the optimization of both source relevance and temporal pertinence of the content allow us to obtain relevant content.

Using keywords and respective popularity from the place model we recur to a news aggregator for searching sources according to the needs of the place model. To make a pre-selection of those sources, anticipate a possible relevance level for a specific source and thus avoiding those sources that are not relevant, we try to know what other people think about each source. This gives us an important marker that can help to improve the system utility, pre-selecting relevant sources and eliminating those sources that no one considers relevant. Important indicators for this measure include the percentage of all Internet users who visit or subscribe a given site or its traffic volume.

At the final of this sub-system, only those sources that are appropriated for the current place context are selected and subsequently, only the most relevant are preserved for the recommender.

C. Scheduler

To take benefit from the potential of digital displays in public spaces, content selection should be adaptive to the dynamic nature of the social environment surrounding the display and it should be flexible enough to react accordingly to new situations in that particular social environment. Furthermore, besides being based in rules or priorities, the content selection should be also based in content relevance models to optimize the relevance of the entire display system.

Content selection is performed taking into consideration public displays scenarios particularities, namely: scheduler should adapt its decisions to the social environment around the display; it needs to balance between the place owner control and the interests of place visitors; and there are some

content that is supposed to be presented more than once, depending on its relevance and the target population. Thus the selection of the next content is realized taking into consideration four important parameters:

a) *Timeliness*: the content relevance depends on its temporal pertinence. This notion of timeliness is of obvious importance in setting the relevance for any type of source, but different sources will handle the effect of time differently. For most sources, the relevance measure should guarantee that the information has not lost its value since publication, but in some cases, a higher relevance may be associated with a particular point in time, e.g. the day of an event, and not necessarily decay as time goes by. More details about timeliness can be found in [12].

b) *Fitness to the place model*: selected content should be appropriate to the social context where it will be presented. Thus, keywords and correspondent popularity from the place model are used to select the next topic to be displayed.

c) *Recommender history*: the scheduler should be aware of the recent and current state of the system. This is important to avoid presenting content that is very similar, although it can be supplied from distinct sources. It is also important for acting as a guide for the display behavior, supporting some behavioral rules defined by the place owner e.g. do not repeat the same item each n scheduling.

d) *Content itself*: because content is provided by distinct sources and may include distinct media content like: text, images, video, etc., it is important to analyze some of its fields. For example, we analyze the length of the text, the content language, the number of links founded in the text and we search for image links. This information is used to influence the content relevance but also to improve the content presentation.

A combination of these parameters, jointly with the scheduler behavior configuration, supports the scheduler decisions of what to present and when to present.

IV. EXPERIMENTAL WORK

To be successful, ubiquitous computing (ubicmp) applications must be designed with their environment and users in mind and evaluated to validate that they do not disrupt the users' natural workflow [13]. Evaluation of ubicmp applications is essential at all design stages. The best designs should include evaluations that involve users in the design process and evaluations should be done repeatedly throughout a series of design iterations.

During the implementation of some modules, complete versions of their prototypes were tested and evaluated in real scenarios, thus validating both the theoretical model and its architecture. The feedback and results we obtained through these experiments, contributed to improve architectural designs as these prototypes were progressively built and also to improve their accuracy.

This section describes the software architecture and two partial experiments, their evaluation goals and their results.

A. Software Infrastructure

Fig. 2 presents the solution scheme that implements the conceptual model defined in previous sections. The four main components of the architecture are: place owner and place visitors preferences specifications; the place model which interprets place seeds and situated keywords and manage the most relevant keywords for the current state of the place; the selector which is responsible for find in the web, relevant sources for the current place; and the scheduler, responsible for selecting the most relevant content to be presented next in the display.

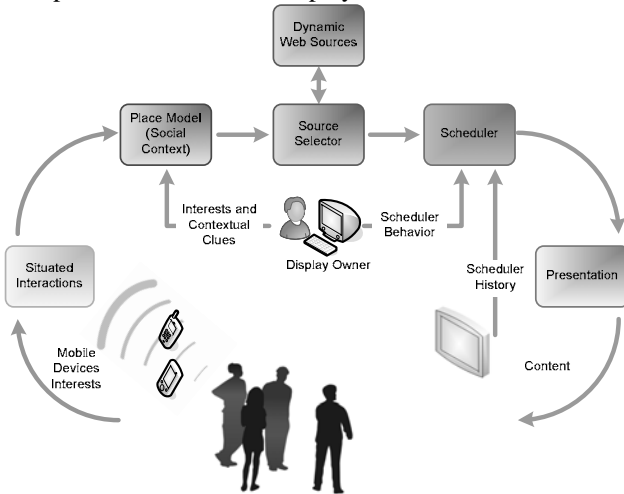


Figure 2. Solution scheme: concretization of the architecture.

B. Experiments and Evaluation Goals

To evaluate the recommender system we have conducted two separate experiments, each one realized in two stages.

The first one was realized to evaluate timeliness models and also to uncover any meaningful user perspectives on timeliness. The second one evaluates the appropriateness and accuracy of the method for obtaining relevant sources based on simple keywords from the place model.

Experiments were realized in real scenarios, out of the controlled environment that is usually founded in lab environments. This is crucial to obtain valid and richness data about users' evaluations. In fact, there are two central issues that make the evaluation process a valid and valuable experiment: evaluate the systems in real scenarios and the involvement of the users in the experiment evaluation.

Next we describe separately each one of these experiments.

1) Timeliness

This experiment was realized to evaluate users' perception on timeliness. It embraced the following set of research objectives:

- a) To understand the key criteria for evaluating the timeliness of content across several types of dynamic sources.
- b) To propose and validate a model for timeliness.
- c) To uncover any elements that may affect people's perception of timeliness.

We proposed two timeliness formulas: one for sources which the relevance measure should guarantee that the information has not lost its value since publication (e.g. news, blogs); other for sources which higher relevance may be associated with a particular point in time (e.g. events).

For 3-weeks the system selects content from a set of 117 dynamic sources of general interest for our target community of people at the University Campus. The system was composed by two displays and it was installed in the hall of the general office of the information systems department. One display showing content from dynamic sources and other, a small touch screen, receiving feedback from users on the timeliness of the content. Results have shown that users' perception on timeliness is well represented by the concept of timeliness as supported by our formulas.

This evaluation with users was complemented with an evaluation of the fairness of the algorithm when choosing between multiple types of dynamic sources. Evaluation on fairness has shown that there are multiple factors that must be considered to ensure a balanced selection among multiple categories or even among the various sources in the same category. This is in part due to their natural dynamic, but still it is an indicator that some of the parameters in the formulas may have to be fine tuned to increase fairness. We also observe that only part of the items was ever displayed. This was a natural consequence of the fact that we had a much higher number of potential data items than time to present them all. Global results also show that the model is able to make comparative calculations of timeliness for different types of dynamic source

2) Keyword-based content generator

The main goal of this experiment was to evaluate if place owners are able to specify their interests in the form of seed tags (simple keywords) and that feeds suggested by the system match the expectations of the seeds' creators. This broad goal embraces to answer to three main questions:

- a) Place Managers are able to successfully specify its preferences using simple keywords?
- b) Place managers consider that their preferences can be described using a small set of keywords?
- c) Place Managers consider that sources provided by the system are adequate to their seeds specifications?

This experiment was realized in two stages. First we ask to 25 participants (researchers and students in different computer science areas) to play the role of the place owner. We asked them to specify their preferences using at maximum four keywords. Based on those keywords the system autonomously selected a set of feeds from Internet sources. The system was also pre-configured with another set of feeds from five distinct topics: generalist, sports, business, technology and culture. Each participant is shown 30 sources (20 derived from specifications and 10 randomly selected from the predefined list of sources). When a feed was presented, there was no indication whatsoever of the selection process or keyword that had been used for selecting that particular feed. For each feed presented, participants should answer to the question "How appropriate is the source considering your seeds specification?" for which they could

give one of the following answers: "Don't know", "Not related at all", "Somehow related" and "Strongly Related". Each participant also fills a simple questionnaire about the simplicity and appropriateness of using keywords-based methods. Results have shown that the majority of all participants specify their preferences easily ($\cong 80\%$) and they obtain appropriate or very appropriate results ($>67\%$). Additionally, 96% of participants refer that their preferences can be specified appropriately using at maximum 6 keywords.

We also realized a two weeks experiment to evaluate how these methods are used and perceived by users in a real scenario. In this experiment to specify place interests we invited four people of the Information Systems Department (DSI) general office to play the place owner role. Each one specified ten keywords and we run this experiment for 2 weeks. The display was installed in the hall of the DSI general office. At the end of the experiment we interviewed each one of them. However, during the interviews several users mentioned situation in which they had witnessed content that they considered to be totally misplaced. For, example, while employment was one of the keywords indicated, they observed that job offers in remote countries was not a relevant content.

Results from both experiments have shown that keywords are normally efficient as content selectors, but they are not always reliable as representatives of the concepts that people had in mind when proposing them. This raises the need of giving to them an interpretation aligned with the nature of place.

V. CONCLUSIONS

In the research carried out within this work we explore different ways to evaluate Ubicomp systems and in particular public displays, to evaluate partial prototypes of our recommender system. Design and implementation processes were realized in a strictly connection with the evaluation process, as is the case of timeliness and also the place description using simple keywords in the place model. The feedback and results we obtained through these experiments contributed to improve architectural designs as these prototypes were progressively built and also to improve their accuracy.

Three main conclusions should be set off. First, the representation of the place owner contributions through simple weighted-keywords proves to be a simple and valid approach. However, their efficiency as seeds for user-generated content may vary greatly. We think we may considerably improve the selection process by framing these keywords in the place context using simple contextual clues. Second, dynamic web sources prove to be an important provider for situated content for public displays. However, it is very important the use of relevance measures to pre-select only those sources that are more relevant ones. Third, timeliness is a central relevance measure and it is essential to improve the relevance of the data items selected for presentation in public displays.

As future work we expect to include in the recommender system information from place visitors situated interactions and also the ability to contextualize these contributions within the place context provided by simple contextual clues associated to the place where the display is installed. We also expect to deploy and to evaluate a fully working prototype of the system.

ACKNOWLEDGMENT

The first author was supported by a Portuguese Foundation for Science and Technology scholarship (SFRH/BD/31292/2006).

REFERENCES

- [1] M. Weiser, "Some Computer Science Issues in Ubiquitous Computing," *Communications of the ACM*, vol. 36, pp. 75-84, 1993.
- [2] E. M. Huang, A. Koster, and J. Borchers, "Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays?," presented at 6th International Conference on Pervasive Computing Sydney, Australia, 2009.
- [3] D. W. McDonald, J. F. McCarthy, S. Soroczak, D. H. Nguyen, and A. M. Rashid, "Proactive Displays: Supporting Awareness in Fluid Social Environments," *ACM Transactions on Computer-Human Interaction*, vol. 14, 2008.
- [4] J. F. McCarthy, D. H. Nguyen, A. M. Rashid, and S. Soroczak, "Proactive Displays & The Experience UbiComp Project," presented at 5th International Conference on Ubiquitous Computing, Seattle, Washington, 2003.
- [5] S. Izadi, H. Brignull, T. Rodden, Y. Rogers, and M. Underwood, "Dynamo: A public interactive surface supporting the cooperative sharing and exchange of media," presented at Symposium on User Interface Software and Technology Vancouver, Canada, 2003.
- [6] D. M. Russell, "Large interactive public displays: Use patterns, support patterns, community patterns," presented at Workshop on Collaboration with Interactive Walls and Tables, 2002.
- [7] T. Payne, E. David, N. R. Jennings, and M. Sharifi, "Auction Mechanisms for Efficient Advertisement Selection on Public Displays," presented at European Conference on Artificial Intelligence, 2006.
- [8] J. Müller, A. Schlottmann, and A. Krüger, "Self-optimizing Digital Signage Advertising," presented at Adjunct Proceedings of Ubicomp, Innsbruck 2007.
- [9] T. Bogers and A. v. d. Bosch, "Comparing and evaluating information retrieval algorithms for news recommendation " presented at ACM conference on Recommender systems, Minneapolis, USA 2007.
- [10] A. S. Das, M. Datar, A. Garg, and S. Rajaram, "Google news personalization: scalable online collaborative filtering," presented at 16th international conference on World Wide Web, Banff, Alberta, Canada 2007.
- [11] R. José, N. Otero, S. Izadi, and R. Harper, "Instant Places: Using Bluetooth for Situated Interaction in Public Displays," in *IEEE Pervasive Computing*, vol. 7, 2008, pp. 52-57.
- [12] F. R. Ribeiro and R. José, "Timeliness for dynamic source selection in situated public displays " presented at 5th International Conference on Web Information Systems and Technologies, Lisbon, 2009.
- [13] S. Consolvo, L. Arnstein, and B. R. Franza, "User Study Techniques in the Design and Evaluation of a Ubicomp Environment," presented at Fourth International Conference on Ubiquitous Computing (Ubicomp2002), 2002.
- [14] F. R. Ribeiro and R. José, "Proactive Scheduling for Situated Displays," in *Workshop on Ambient Intelligence Technologies and Applications*. Guimarães, 2007.

This article was featured in

computing **now**

ACCESS | DISCOVER | ENGAGE

For access to more content from the IEEE Computer Society,
see computingnow.computer.org.



IEEE

IEEE  **computer society**

Top articles, podcasts, and more.



computingnow.computer.org