

Smart content selection for public displays in ambient intelligence environments

Fernando Reinaldo Ribeiro

Rui José

Abstract A public display that is able to present the right information at the right time is a very compelling concept. However, realising or even approaching this ability to autonomously select appropriate content based on some interpretation of the surrounding social context represents a major challenge. This article provides an overview of the key challenges involved and an exploration of some of the main alternatives available. It also describes a novel place-based content adaptation system that autonomously selects from web sources the content deemed more relevant according to a dynamic place model. This model is based on a tag cloud that combines content suggestions expressed by multiple place visitors with those expressed by the place owner. Evaluation results have shown that a place tag cloud can provide a valuable approach to this issue and that people recognize and understand the sensitivity of the system to their demands.

Keywords: ambient intelligence, public displays, context-awareness, recommender systems, human-centred, ubiquitous computing.

1 Introduction

In Ambient Intelligence scenarios, the environment can perceive and react to people, sense on-going Human activities and proactively respond to them. Public digital displays have always been part of this vision (the “boards” (Weiser, 1993)), and their increasingly ubiquitous presence in our socio-digital landscape has been opening new opportunities for their use as important building blocks for many types of Ambient Intelligence multimedia environments. However, most public displays today are mere distribution points for pre-defined and centrally created content. They assume passive users and they are completely unaware of the Human activities taking place at that same location. Consequently, they are not very valued by their potential users and their content is often perceived as too institutional or dull (Huang, Koster, & Borchers, 2009).

The solution may be found in their ability to dynamically integrate content from the web and select sources according to their relevance to the social context around the display. The number and diversity of content sources on the Internet offers the potential to guarantee a continuously updated stream of relevant content for the displays. The potential is so vast that we can safely say that content would no longer be a scarce resource and that proper selection would indeed become the key problem. Mobile technology, on the other hand, may offer the potential to implement multiple forms of automated personalization, or adaptation, of the displays. The possibility to collect preference expressions from people near the display could enable content that is tailored to the preferences or goals of the individuals or groups visiting the place where the display is set, considerably improving the user experience and enriching the place characterization with the preferences and goals of visitors. 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 the people in a particular place. The problem of place-based

content adaptation in public displays may thus be formulated as follows: Given a set of preferences implicitly or explicitly expressed by the people around a display, how can they be combined with the characteristics of the place to enable the system to select from web sources the most appropriate content to be displayed in that context.

This paper reports on the final results of a research program aiming to uncover content adaptation strategies for public displays based on the preferences of nearby people. It provides an overview of the key challenges involved and an exploration of some of the main alternatives available. As part of this work, we have developed a public display system that autonomously selects from web sources the content deemed more relevant according to a dynamic place model that is sensitive to the people around the display. People can express their content preferences by specifying tags in their Bluetooth device name, as described in (José, Otero, Izadi, & Harper, 2008). These tags, together with seed tags defined by a place owner constitute a tag cloud that is continuously reflecting the social setting around the display, being sensitive to immediate indications of interest and providing a balanced combination between content suggestions expressed by multiple place visitors and those expressed by the place owner. This tag cloud is then used as the basis for an adaption process that addresses the specific requirements raised by content selection in public displays. The results obtained with the evaluation studies show that place visitors recognize the sensitivity of the system to their demands and that a place tag cloud can provide an important element for the interpretation of place and for combining the dynamic set of interests expressed by multiple people. Overall, these studies provide a relevant contribution towards understanding the implications of adaptation approaches for public displays and more generally to the discussion in the community regarding adaptive behaviour in public spaces.

2 Challenges in adaptive content for public displays

The idea of a public display that is able to present the right information to the right users, at the right time and in the right way is obviously very compelling. However, realising it, or even approaching it in some way, is extremely challenging because of the complex issues involved, such as obtaining information about preferences in a non-obtrusive and privacy-preserving way, the need to combine the various preferences expressions of the people in that place, and the complexities involved with making meaningful inferences about relevance based on the information available. Instead of a pre-defined schedule, the system would work based on policies that determine the high-level behaviour of the system, but not exactly what it will do. In the end, the effective behaviour exhibited by the system will result from the combination between those policies and the *stimuli* received from the environment. When considering how to approach this goal, we have explored two major alternatives: modelling this problem as a context-awareness problem and modelling it as a recommendation problem.

2.1 Adaptive scheduling as context-awareness

Our first approach to support adaptive scheduling in public displays was to model this as a context-awareness problem. A context-aware scheduler should be able to access data about the display context and make scheduling decisions that reflect rational choices about what is most appropriate for each particular context. This requires some type of knowledge that associates context states with specific content. We have explored three alternatives for creating this association: user-generated rules, supervised learning and unsupervised learning.

A first possibility is to embed into the scheduling process a number of behaviour rules that represent our empirical knowledge of 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 scheduled for presentation when the local temperature is above a certain threshold, or that another item should not be shown twice in the presence of any particular Bluetooth device, or that yet another item should only be shown when there is no one being detected in the immediate physical proximity of the display. With this type of context-based rules content can be selected according to the continuously changing context around the display. We developed a display system in which content was determined by a set of context-based rules. A JESS Rule Engine was used to evaluate those rules and increase, decrease or eliminate the utility associate with the various content items. The rules included references to context variables, such as nearby devices, time of day, content characteristics and noise level. The main observation from this early study was the recognition that specifying meaningful rules that can apply to generic usage situations can be extremely difficult. Even though it seems intuitive to think that certain situations can affect what is the most appropriate content to be presented in a particular context, it is not trivial to explicitly translate, based only on empirical knowledge, subtle interpretations of context into formal and generic specifications that will rule the display behaviour in multiple real-world contexts. Smartness seems to be a rather vague concept that for most cases does not map directly into just one predefined type of reaction. Considering these challenges, a clear alternative towards context-aware adaption in public displays is to allow the display system to learn a particular notion of smart behaviour, by training the system to generate new domain knowledge.

In a supervised learning process, there would be a training stage, in which a set of training cases of the expected behaviour would be generated, an inference stage, 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 unclear, however, how much of the generated knowledge would be generic enough to be applied to multiple displays. Each display will have its own context and content items and, moreover, these will be evolving with time. Considering these limitations we have not explored this particular approach.

In a process of unsupervised learning, the scheduler would be able to learn with experience and adapt to the specific characteristics of its environment. This would require a generic feedback mechanism whereby its users would be invited to express their opinion on the scheduling decisions, or more specifically about the relevance of what is being shown. This learning process could lead to two complimentary types of result: information about the global popularity of individual items; and information about the effect of context in content popularity. To explore the potential of this approach, we conducted a user study to assess what type of context variables could possibly affect people's perception about content relevance. For that purpose, we deployed the experimental setting represented in Fig. 1.

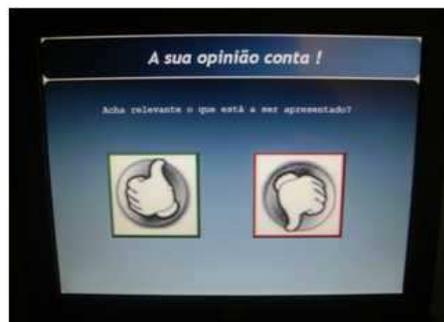


Fig. 1. *a)* Main display on top and touch-screen underneath; *b)* Touch-screen for feedback

A public display was used to present general purpose information, such as news (headlines, technology, and culture), advices, weather, department website, university website and a slide show with photos related to the University. The scheduler was not context-aware, as it simply cycled through the content throughout the day. A second, touch-based display next to the first one, allowed people to provide positive or negative feedback about the relevance of what was being shown on the main display. Every time someone expressed an opinion, we stored information about the content item, the feedback given, and the current status of the context variables, such as time and the list of currently present Bluetooth devices. The results have shown that particular content items were systematically rated more relevant, regardless of any context changes. From all the tests we have made, we have found no statistical significance whatsoever in any of the associations between the relevance of a particular item and any of the context variables. People seemed to associate the relevance of content much more with the place in which it was being presented rather than with any of the forms of context being sensed.

These first experiments have revealed important limitations in the overall concept of context-aware adaptation in public displays. The difficulties in expressing or inferring meaningful associations between context variables and content items mean that the scope of situation in which those approaches can be effective will not be as broad as is often suggested. While we do not claim to have extensively explored all the alternatives and technological approaches, it seemed clear that context-awareness would only be efficient in cases where we have a very direct association between a particular context state and a specific system reaction, normally something in the form of a trigger. These findings for context-aware displays are in line with more generic findings for context-awareness e.g. context-aware systems should promote a tight integration between sensing and action instead of decomposing them as separate parts of the same process (Leahu, Sengers, & Mateas, 2008) the relevance of the sensor readings corresponding to context variables cannot be established a priori outside the specific situation of use in which it is being generated (Dourish, 2004). These conclusions have strongly influenced our work and lead us to decide to place our efforts on alternative approaches. Instead of focusing on sensing, modelling and inferring context, we decided to focus on the situated meaning of people's interactions with the system and also on data about content items, as the main drivers for the adaptive behaviour of system.

2.2 Adaptive scheduling as a recommender system

The specific scenario of recommending content from web sources for presentation on a public display raises new challenges that break some of the assumptions we often find in recommendation algorithms and that may limit their applicability. This section summarises the main differences between both problem domains.

Limited set of preference indications

Unlike other media, attention and engagement are far from being a given with public displays. In most cases, the system is autonomously selecting what to presented next and people are very limited in their ability to influence the display decisions. This happens, not just because of technical limitations resulting from the lack of a mouse and keyboard, but essentially because the display is public and not under the control of any single user. Most public displays are thus conceived with the single goal of being seen by the people in their vicinity, without considering any other forms of interaction. Even in the few cases

in which there is some interaction support, the richness of those interactions, and consequently, the richness of the digital footprints they generate, is much more limited than in the more traditional scenarios of recommender systems, such as the web. Public displays can thus be used without generating any relevant information about how they are used, much less about the users' interests or preferences. Without such information, it is not feasible to implement recommendation techniques that depend very heavily on implicit or explicit expressions of interest.

Profiles for Place-based adaptation

While most recommendation systems target individual recommendations, public displays must be designed for shared and communal use in public and semi-public settings. Therefore, instead of a user profile, recommendations for public display should be based on some sort of place profile that combines the preferences of the person managing the display, who we call the place owner, with the preferences of the multiple people that may be in the vicinity of the display.

A place owner who installs a public display will have specific expectations regarding the way in which the display is going to contribute to the creation of a particular concept of place. An effective profile model must acknowledge this role of public displays and provide the place owner with some control over the nature and scope of the recommendations. If the selected content is not aligned with the place values, practices or commercial strategies, its public presentation may become a source of embarrassment.

Place visitors, on the other hand, are the main consumers of the content presented on the display and for that reason they must also play an active role in guiding the display behaviour. Since public displays will typically have multiple simultaneous users, the adaptation process will need to consider the best strategy for dealing with the potentially very varied interests expressed by those people. This generates a trade-off between the selection based on a profile combining the multiple interests of the multiple persons present and the selection based on the use of each individual profile, one at the time (Schilit et al.). The first is a balanced approach, but faces the risk of not really matching anyone's specific interests. The second approach can be targeted for each individual, but it raises additional privacy issues and may conflict with the idea of public displays as a place-making tool.

Select, not just recommend

Most recommender systems assume some type of collaborative process with users. The system may suggest multiple data items from various sources and present them in the form of short summaries with links for further details, but the user is then expected to assess the multiple relevance cues provided and select which content may be of interest. For example, when browsing a video sharing web site, new video suggestions are normally presented after a video is finished. This gives users the possibility to explore the video collection in a more serendipitous way, but they are also free to ignore the suggestion and follow their own path. In Internet radio systems, the system autonomously selects the next music, but the user can easily override selections by skipping to a next song. There is, therefore, an implicit assumption that the user will somehow be in control of the selection process and that the system responsibility is mainly to facilitate that process.

In a public display, none of these possibilities is naturally available. Since there is normally no intermediate stage between content being suggested and being presented, the selections of the system will be shown without necessarily being evaluated by a Human, and possibly without any immediate mechanisms for skipping it. Furthermore, given that people will not normally have the possibility to control presentation, content will normally

be shown in its full extent. This raises considerably higher risks of presenting irrelevant or even inappropriate content and substantially raises the responsibility of the system in being able to only recommend appropriate items.

Dynamic and open content sources

A significant part of the research related to both, public displays and recommendation systems, assumes that there is a predefined list of content or a predefined list of content categories (Churchill, Nelson, Denoue, Helfman, & Murphy, 2006; Greenberg & Roudind, 2001; Interrelativity Inc., 2005; Izadi, Brignull, Rodden, Rogers, & Underwood, 2003; J. F. McCarthy, 2000; Joseph F. McCarthy, Costa, & Liongosari, 2001; Joseph F. McCarthy, McDonald, Soroczak, Nguyen, & Rashid, 2004; Russell & Sue, 2002) from which content is selected. This strategy has the advantage of allowing a strict control over the type of content that can be presented on the display, but it would fail to address our initial motivation of benefiting from the wealth of content and information sources on the web. To really take advantage of the potential of web information, the display system should allow display users to freely express their interests and be able to dynamically discover and select from the Internet the most adequate sources.

The focus on dynamic content from web sources also means that the relevance of the respective content is likely to face considerable oscillations. The same source may at a given moment have rich and recently updated information and at some other moment strongly deprecated or even non existing information. The notion of relevant source may thus change very quickly, not just because of changes in the display context, but also because of changes in the content itself. As a consequence, the adaptation processes at the various selection levels must all be very dynamic and frequently re-evaluate their selection decisions. Moreover, previous feedback on a particular source may itself become deprecated very quickly, given that the respective content may have changed considerably.

Presentation cycles

In a recommender system, content is presented once and after that the system assumes the user is no longer interested in the same content. In a public display, the same content may be presented multiple times because it will most likely be seen by different people at different times. A recommendation system for public displays needs to take into account the existing of presentation cycles in which it becomes acceptable to show the same content again, if it remains relevant. The calculation of relevance may thus have to balance the inherent relevance of a content item with the effect on that relevance of previous presentations of that same content item.

3 Related work

This work combines contributions from multiple areas. In this section, we present an analysis of previous work structured around the topics of adaptive scheduling in public displays, recommender systems and use of tag clouds for interaction and profile representation.

Scheduling in situated displays

For most system, the scheduling is a based on a fixed schedule that cycles through predefined content, but several adaptive scheduling alternatives have been explored that introduce sensibility of the display to some type of external variable. Proactive displays (Joseph F. McCarthy et al., 2004; McDonald, McCarthy, Soroczak, Nguyen, & Rashid,

2008) select content that is scheduled on-the-fly according to the interests of users within the direct vicinity of the display. It recognizes specific presences and display information from associated profiles. Similarly, Groupcast (Joseph F. McCarthy et al., 2001) uses identification and user profiles to present information about the interests of people in the vicinity of the display. Context-aware informative display (Zhu, Zhang, Zhang, & Lim, 2007) and context sensitive public display for adaptive multi-user information visualization (Morales-Aranda & Mayora-Ibarra, 2007) also aim to provide people with relevant content in an opportune and personalized way. Dynamo (Izadi et al., 2003) or BlueBoard (Russell, 2003) are examples of systems that give users direct control of the display and thus content selection is directly handled by users rather than by the system. The BlueScreen (Payne, David, Jennings, & Sharifi, 2006) selects and displays adverts in response to users detected in the audience. It utilizes Bluetooth-enable devices as proxies for identifying users and utilizes history information of past users' exposure to certain sets of adverts. Advertisements are preferentially shown to those users that have not seen them yet. Muller (Müller, Kruger, & Kuflik, 2007) describes a mechanism to adapt advertisements on digital signage to the interests of the audience. He proposes a system that uses a naïve Bayes classifier to estimate the probability that a user is interested in a certain advertisement. It uses adverts keywords, users' history, time, location and voucher collection information as 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, mostly based on individual profiles. In our work, we assume that we have no a priori knowledge about users' profiles, and use a place specification as the basis for autonomous content selection. Furthermore, our approach is not specific to any particular type of content. Even though in this paper we are only addressing information feeds, the same specification could be used for selecting videos, photos, advertisements, or any other type of content.

Recommender systems

Regarding previous work in recommender systems and information retrieval, we tried as much as possible to leverage on existing algorithms and tools. Still, as described in Section 2, there are several specificities associated with this particular problem domain that limit the applicability of existing techniques. Approaches to recommender systems deal mainly with two types of entities: users and items. The user entity is normally a user profile that is based on either manual user input of preferences or automatic user modelling i.e. deriving user preferences and providing recommendations on the basis of user's history of content consumption. The item entity is usually characterized with a set of metadata that is supplied by the source, but this information can also be extended with additional information that is inserted by users. These entities are the basis of all recommendation techniques that normally depend on extensive data about both. However, the specificities of how they are used differ between the three main recommendation techniques in use today: content-based recommendations, collaborative recommendations and hybrid approaches.

In content-based recommendation, the system suggests to the user the items that best fit the user profile. A set of attributes that characterize each item is used to determine appropriateness of the item for recommendation purposes. The user will be recommended items that are similar to the ones preferred in the past. A key limitation of these techniques is that the system cannot recommend items that are different from anything the user has seen before. This is particularly limiting in our case because the ability to find new content that is totally unrelated with what has already been presented is an important requirement.

Collaborative techniques, in their common form, are also difficult to apply. They make recommendations based on information of people with similar tastes and preferences, trying to predict the utility of items for a particular user based on the items previously rated by other users. Their main limitation is not being very good in dealing with new content items and also with frequently updated sources (Das, Datar, Garg, & Rajaram, 2007), making them unsuitable for our scenarios of selecting content feeds in a timely way. Moreover, the focus on a single display, the reduced number of users in our system and also the potentially very low rating density raise additional challenges to the application of collaborative recommendation techniques.

González et. al (González et al., 2006) present a system that extends traditional approaches to recommender systems. They analysed cross-disciplinary trends from the users' affective factors perspective in the next generation of ambient recommender systems. They combine a model of the user's emotional information with intelligent agents and machine learning to provide relevant recommendations in everyday life. This work also builds on previous work in recommendation systems and retrieval models for feed search (Arguello, Elsas, Callan, & Carbonell, 2008; Bihun et al., 2007; Seo & Croft, 2007).

Keyword-based interaction

Previous work has also reported on the use of tag clouds to represent realities outside their original context, including personal profiles, social dynamics in meetings or trends in political discourse. Steinbock et.al studied the use of wearable tag clouds in face-to-face interaction (Steinbock, Pea, & Reeves, 2007). Within the context of an academic meeting, participants were given a large badge with a tag cloud of the most common words in their published documents. This was expected to represent a synopsis of the respective interests and facilitate interaction between participants. McNaught and Lam (McNaught & Lam, 2010) explored tag clouds to analyse the spoken and written responses of informants in focus groups transcripts. The tag clouds facilitate the study of the social dynamics in those groups, but this analysis is only conducted at a later stage and provides no feedback on the social interaction as it unfolds. Viégas and Wattenberg report on some of the uses of IBM site Many Eyes, where people upload and visualize data in a variety of ways (Viegas & Wattenberg, 2006). An example includes the use of a tag cloud, created from the set of blogs a person normally reads, as a representation of a personal profile. Tagline Generator (Mehta, 2006) supports the generation of chronological tag clouds from time-based text data sources. This work is an example of another system that deals with the time dimension. In these tag clouds, the colour in the words is associated with usage variations. Words whose usage is increasing will be brightened, while words whose usage is decreasing will be fading away. The use of tag clouds for content recommendation has been described by Pessemier et.al (Pessemier, Deryckere, & Martens, 2009). Tag clouds are generated from user ratings to create a form of personal profile. These tag clouds are then used to recommend movies to that person. In our work, we also suggest the use of tag clouds for recommendation purposes, but in our case this corresponds to a place or situation profile, and not to the profile of a single individual. This related work demonstrates how tag cloud can indeed support many roles outside their original context. However, the use of tag clouds as a situated representation of place that drives the content selection on a public displays remains to the best of our knowledge a novel approach.

4 Adaptive place-based content scheduling

As part of our work on this topic, we have developed an adaptive content scheduler for public displays that is able to select the information feeds deemed more relevant for the current social setting around the display. The respective design has evolved throughout the various prototypes and its final architecture is represented in Fig.2.

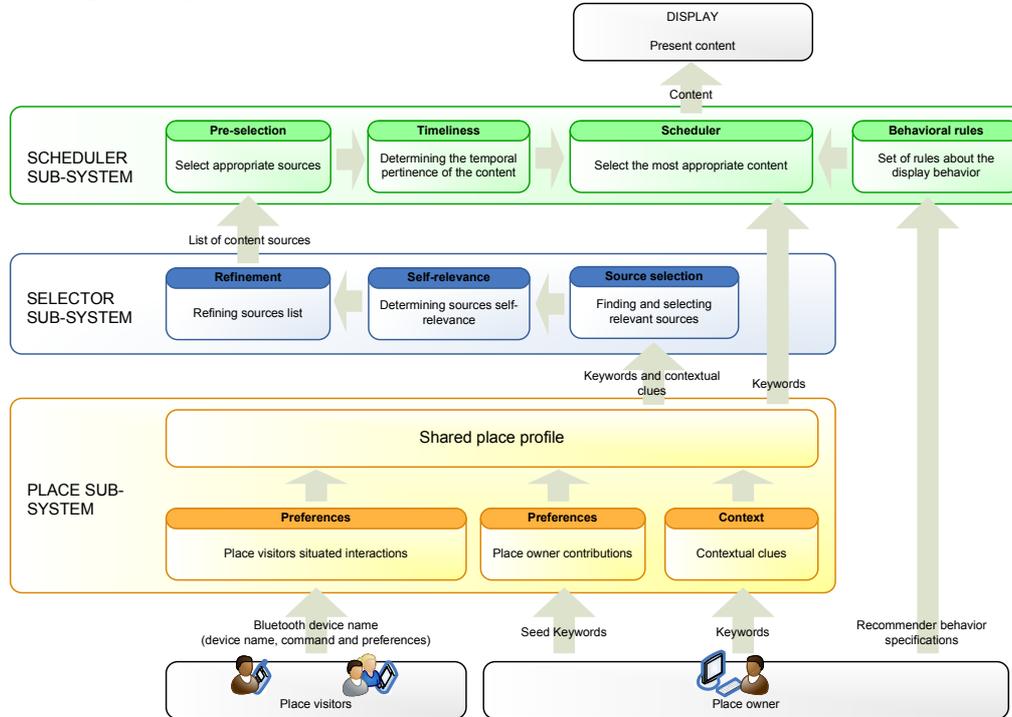


Fig. 2. System Architecture.

There are three main sub-systems in this architecture: a place sub-system representing a dynamic and evolving view of place that combines the place owner specifications with the contributions made by place visitors; a selector sub-system that takes keywords from tag cloud in the place sub-system and retrieves relevant content for each of those tags using dynamic web sources as content providers; and a scheduler sub-system that considers the available content as well as the weight and presence level of the represented tags to select which content is going to be shown next. We will now describe the operation of each of these sub-systems in more detail.

Place sub-system

Place is a very broad concept, involving a physical setting, situated practices and a fluid interpretation of the surrounding social context. A place model should thus combine the place-making role of a place owner with the ability to allow the multiple place visitors to also exert some influence on the characteristics of place.

The model used in this system to represent place is based on the use of a place tag cloud as a shared and evolving view of the expressions of interest made by the place owner and the people around the display. Tag clouds support comprehension by offering a representation of the topics associated with a place and they support navigation by offering an aggregate interaction mechanism, whereby keywords in the tag cloud provide an immediately available and dynamically evolving list of interaction suggestions. These are valuable features, even if the reality being represented is no longer a set of documents or resources, but a stream of interactions occurring around a public display. The tag cloud

becomes the unifying concept that associates keyword inputs with dynamic behaviours, functioning as the key driver for the selection of place-based content.

The place tag cloud is first created by the place owner. Place-making parameters allow a place owner to provide additional characterization and specify adaptation boundaries for a tag cloud. Even though the tag cloud is expected to emerge from interaction, these parameters provide a way for aligning the display behaviour with the general expectations of appropriateness of the display owner and its place-making objectives. The main part of these specifications is a set of keyword lists that enable some control of the tags in the tag cloud. The first is a blocked words list. This list may include words that are very common in a given language, but not very meaningful in a tag cloud. Excluding these words is already a common procedure in traditional tag clouds, but in this case we may also want to use this to prevent abusive or offensive keywords from making it into the tag cloud. We also have a list of seed keywords that serve to initialise the tag cloud and maintain a number of place keywords when there are not enough keywords being generated. Seed tags are defined with a minimum popularity value that determines how visible they remain when other tags begin to emerge. A *SeedsOnly* parameter can be used to determine that a particular tag cloud will only accept seed words. This works as a white list that restrains the accepted words to those on the list and may be useful to promote aggregation around thematic tag clouds, e.g. a tag cloud with sport teams, emoticon symbols, or music styles. Finally, there is also a list of contextual keywords that can be used to provide additional context to the words in the tag cloud. For example, Sports and Football could be added to a tag cloud representing football teams. This is particularly important if the tag cloud is to be used for selecting web content.

The tag cloud will then evolve with the continuous stream of words being generated by various types of implicit and explicit interactions, such as Bluetooth names, Obex exchanges or SMS/MMS messages. In our prototype, place visitors could publish their own tags into the tag cloud, by including *tag* commands in their Bluetooth device name, as described in (José et al., 2008). Tags that have just been generated have a special meaning because they correspond to the presence or interaction events created by people who are now in the vicinity of the display. The presence of tags represents an additional dimension that is not normally included in traditional tag clouds, but may be key to interpret the immediate relevance of those tags in content selection, possibly favouring tags that are currently present, albeit less popular, instead of popular tags that no one is currently generating. The tag cloud resulting from combining the initial specification by the place-owner with the continuous stream of words originating from presence and interaction provides the evolving and dynamic representation of place that will be used as the basis for content selection.

Selector sub-system

The selector sub-system is responsible for autonomously finding and selecting relevant content sources from the Internet. To fully benefit from the wealth of sources on the Internet, the selector should be able to consider the generic relevance of a source, in the sense of popularity, and also its timeliness, in the sense of how up to date is the respective information.

Using keywords from the place tag cloud we recur to a feeds aggregator for searching sources according to the needs of the place model. The result is a large set of sources without any relevance criterion. We then apply a relevance algorithm to promote sources that rank higher in popularity. This popularity is independent of the usage context where the content is consumed and is based on generic measurements such as the percentage of all Internet users who visit a given site or the traffic to the site. We use Alexa search engine to obtain the traffic rank, a measure obtained from Alexa Toolbar users. An API

provides information about a specific Uri, namely their rank, their traffic and related sites. This information is country specific, and so we also consider the country of origin as an additional input. Another important measure of feed relevance is the number of users subscribed to the feed. Higher values of subscriptions denote higher feed interest. Through the Newsgator API, we get the number of users that subscribe each of the sources. Because both traffic rank and feed subscriptions have distinct numeric domains they need to be normalized. We do that by defining acceptance thresholds for both measures. Sigmoid functions are used for characterizing acceptance intervals and provide a smooth interpolation between the limits of those intervals. Both traffic rank and number of subscribers are thus combined in a single function that determines the most relevant sources for each of the tags associated with a place.

Scheduler sub-system

The scheduler sub-system decides which content to present next on the display. This decision involves two steps. The first is to select which of the tags in the place tag cloud will be used next for content presentation. For each tag, there is information about the respective popularity and whether or not it is currently present, i.e. someone who is now around the display is announcing that tag. Additionally, there is also a list of recently used tags. When selected for presentation, a tag is then placed on this list for a number of iterations to prevent it from being successively selected. The algorithm selects the most popular tag among those who are present and not in the waiting queue, or if none is present, the most popular tag not in the waiting queue.

The second step is the selection of which of the content feeds associated with the selected tag will be displayed next. This decision is based on a multi-criteria utility function that considers three parameters: timeliness, content structure and scheduling history.

Timeliness represents the measure of the temporal pertinence of content and will be described in more detail in the evaluation section.

Content structure is an empirical measure of the appropriateness of the content in the feeds for presentation on a public display. This is not related with the semantics or the quality of content, but only with the fact that, in certain cases, the structure of content can make it inadequate for presentation in a public display, either because it is link-intensive, because the text is too long, or for some other similar reason. Because content is provided by distinct sources and may include distinct media like: text, images, video, etc., it is important to analyse some of its fields. For example, we analyse the length of the text, the content language, the number of links found in the text or the number of image links in the feeds.

Scheduling history considers the previous presentation of similar content. The scheduler should be aware of the recently presented content and avoid presenting content that is very similar, even if supplied from distinct sources. This is very common with news sources, which may be showing very similar headlines.

A combination of these parameters, jointly with the scheduler behaviour configuration, supports the scheduler decisions of what to present at each moment.

5 Evaluation and Discussion

The evaluation of this adaptive content system comprised a set of smaller user studies conducted throughout the project development and a final public deployment to assess its overall operation in a real-world scenario. The initial studies aimed to evaluate specific system features, more specifically the timeliness model and the place specification. In this

section, we will first present a short summary of the key contributions of each of those studies to provide the reader with a global view of the work. The feedback and results we obtained through those experiments, contributed to improve architectural designs as these prototypes were progressively built. This section ends with the description of the final evaluation of the system, which was conducted in a real word scenario and generated rich data about users' perceptions of the adaptation behaviour.

5.1 Assessment of timeliness model

Our first study has aimed to understand the key criteria for evaluating the timeliness of content across several types of dynamic sources (Ribeiro & José, 2009b). Timeliness is a measure of the time-based relevance of a content item and is of obvious importance in setting the relevance for any type of source. Different sources will suffer the effect of time differently, and thus we have used two distinct models to calculate this measure: one for publication-based sources, where the timeliness measure was mainly based on how recently they had been updated (e.g. news, blogs); and another for sources in which timeliness was associated with a particular point in time (e.g. events). Each of these formulas could then be tailored with parameters to make it more suitable to the characteristics of particular sources. The experimental setting comprised a public display showing the selected content, and a second display, a small touch screen, where people could give feedback about the timeliness of the content being presented on the main display. During a 3 weeks experiment, the system selected content from a set of 117 dynamic sources of general interest and collected the feedback from users. This evaluation with users was complemented with a log analysis to evaluate on the fairness of the algorithm when choosing between multiple types of dynamic sources.

Results from the user studies have shown that the users' perception of timeliness was properly represented by the concept of timeliness as supported by our formulas. However, the log analysis has revealed the existence of multiple factors that must be considered to ensure a fair selection across the various content categories or even among the various sources in the same category. Additional research would be needed to calibrate the model to support comparative timeliness evaluations between different types of content source.

5.2 Assessment of the place model

A second study was conducted to evaluate if place owners were able to successfully specify their preferences using a small set of keywords and if they considered that the feeds suggested by the system matched their specifications (Ribeiro & José, 2009a, 2009b). This would provide an important evaluation of the appropriateness and accuracy of the method for obtaining relevant sources based on simple keywords from the place model. This experiment was realized in two stages. First we asked participants (researchers and students in our University) to play the role of place owners, and specify their preferences using a maximum of four keywords. Users would then seat in front of a display in the laboratory to evaluate to what extent the content presented by that display would match their expectations when specifying the keywords. The display system would then start presenting content based on those keywords by autonomously selecting feeds from the Internet. The system was also pre-configured with another set of feeds from five distinct topics: generalist, sports, business, technology and culture. Each participant was then shown 30 sources (20 derived from the 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. At the end, each participant filled a simple questionnaire about the simplicity and appropriateness of using keywords-based methods. Table 1 presents a summary of the evaluations.

Table 1. Evaluation of relevance for keyword-based selection.

	Total sources	Don't Know	Nothing Related	Something Related	Very Related
Sources from keywords	493	27 (5,5%)	154 (31,2%)	122 (24,8%)	190 (38,5%)
Sources from predefined list	233	5 (2,1%)	147 (63,1%)	31 (13,3%)	50 (21,5%)
Sources from predefined list (but related to specified keywords)	17	0 (0,0%)	2 (11,8%)	0 (0,0%)	15 (88,2%)
Sources from predefined list (but not related to specified keywords)	84	1 (1,2%)	73 (86,9%)	8 (9,5%)	2 (2,4%)

Results have shown that the majority of all participants ($\cong 80\%$) specified their preferences easily and 96% of them refer that their preferences can be specified appropriately using at maximum 6 keywords.

With respect to the appropriateness of the content to the specified keywords, they obtain appropriate or very appropriate results in more than 63% of presented contents that are obtained using keywords. In contrast only 34,8% of content obtained from predefined list of sources were evaluated as appropriate or very appropriate. It is important to refer that sources from the predefined list were chosen considering generalist interests. In some cases, these would match the users’ own interests and would therefore generate positive evaluations. This is the case of the third line of the table 1, 88,2% of sources were evaluated as very relevant. However, when we analyse results related to the sources from predefined list there are also a set of sources that are not related to the interests specified by users and in this case 86,9% of sources were evaluated as nothing related.

This has confirmed that the method for obtained relevant sources based on simple keywords is normally efficient, but it has also shown that when users interests match the main topics (categories) of sources from a predefined list (that should be carefully selected according to the nature of the place where the display is situated and generalist interests) the content from these sources could be very relevant. This shows that the combination of the two approaches may represent a promising approach.

As part of these same evaluation goals, we also conducted a secondary experiment in which we invited four people from the general office at the Department to play the role of place owners for a public display installed in a public place outside that office. Each participant started by specifying ten keywords. We then collected and combined those keywords to form a static tag cloud that would drive content selection throughout a two weeks experiment. Participants were given a diary, and asked to record their impressions on the observed content. At the end, we conducted interviews with participants. During the interviews several users mentioned situations 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 representations of the concepts that people had in mind when proposing them. This need to align the tag interpretations with the nature of place has lead to the introduction of place keywords that are never used on their own, but complement words in the tag cloud to give them additional context for the process of feed selection.

5.3 Public evaluation

The final system evaluation was based on the public deployment of a fully functional implementation of our adaptive content scheduler and its global objective was to evaluate under realistic conditions the overall operation of the system. There are many different goals that may be at stake when evaluating recommendation systems. While certain evaluations may focus on the recommendations accuracy to predict ratings, others may focus on how often the system leads its users to wrong choices, the degree to which the recommendations cover the entire set of items or the ability of recommenders to explain their recommendations to users. In our evaluation, we have focused mainly on user perception about the system behaviour rather than on a specific notion of accuracy. It would always be subjective to set any comparison metrics for assessing the appropriateness of content to the social environment around the display. Therefore, we chose to focus on how users perceive the behaviour of the system within their own conceptual frame of the surrounding context. More specifically, we wanted to address the following evaluation goals:

- the appropriateness of the place tag cloud as a shared and public model for the social setting around the display;
- the viability of tag indications as a simple interaction mechanism;
- the recognition by place visitors of the sensitivity of the system to their demands;
- and the overall perception that the content displayed is relevant for the place as expressed by the tag cloud.

The experimental setting comprised a large public display at the main entrance of the Information Systems Department at University of Minho. This is mostly a transient place where students, teachers and staff walk by on their way to classes and labs, occasionally forming small groups. People passing by could publish their own interests through special commands in their Bluetooth device name. The combination of multiple contributions from visitors was expected to provide a seamless sequence of social settings that would drive the display content. The display content was organised around three panels as shown in Fig. 3a.

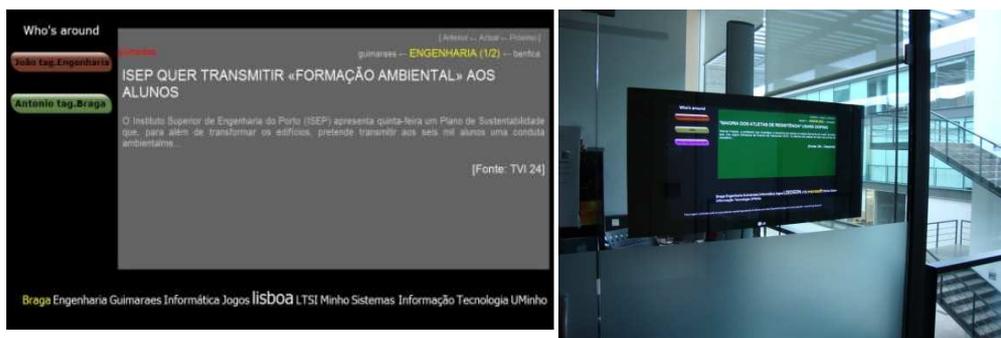


Fig. 3. a) Public display screenshot. b) Prototype image.

The left panel presents Bluetooth device names that are currently detected in the place. This is the “*who is around*” panel and essentially serves as a first level of presence recognition. The bottom panel represents the place tag cloud that evolves with the tags that people can add to their Bluetooth device name. A tag is added to the place tag cloud, if new, and it will be marked as being present, with its popularity steadily increasing while that presence remains. The size of the tags is proportional to their popularity in the place and keywords that are currently present are shown in yellow as an additional level of presence recognition and interaction feedback. This tag cloud was seeded with 20 words representing topics related with Informatics and Engineering as well as keywords

associated with the town and region. The main panel presents the content selected by the system according to the current place profile. For this experiment the display was configured to favour Portuguese content, to avoid presenting the same item every 10 presentations and to avoid presenting items from the same source every 5 items. The scheduler considered both the weight and current presence level of the represented tags to select which content it would show next. The top section of the content area shows the sequence of keywords being used for content selection, more specifically, the previous, the current and the next tags being used, thus facilitating the interpretation of the system behaviour by users.

During the 3 weeks of the experiment we collected usage logs and Bluetooth data corresponding to presence and interaction events. To evaluate people’s perception about the responsiveness of the systems to indications of interest and the relevance of the content presented, we also conducted a total of 15 structured interviews with people who had previously tried to use the system. These people were asked to answer 6 groups of questions related with their experience. For each group there was also an open question where users were able to make their comments about the topic.

5.2.1 Results and discussion

Table 2 summarises the usage statistics observed during the 3 weeks of the experiment between 9am and 8pm of the working days.

Table 2. Usage data from final evaluation

Distinct Bluetooth device names detected	349
Bluetooth devices detected	308
Total explicit interactions (tag commands across multiple sessions)	73
Distinct explicit interactions (unique tag commands)	31
Distinct Bluetooth devices with explicit interactions	23
Number of scheduled items	21616
Number of items selected from interactions originated tags	5567
Number of items selected from tags while those tags were present	360
Number of selected sources	128

The numbers show that there were 23 unique devices originating tags to the display. As a result 1,67% (360) of the scheduled items were presented as a direct reaction to tag publications while the respective user was present. This reflects the relatively limited number of people interacting with the system and above all the transient nature of the place, with most people just stopping for a very short interaction and then moving away. We can also observe that 25,7% of all schedule decisions have originated from users’ interactions (with the remaining 74,3% being based on seed tags). This number is particularly interesting because it reflects the influence of the aggregate of interactions in the display behaviour even after the people making the interactions are gone. This number is strongly affected by the interaction patterns, but also by key behaviour parameters, such as the decay level in the popularity of tags or the minimum popularity defined for seed tags. A quick decay will make the system more responsive to new arriving tags, but it will also fall-back more quickly to the seed tags. A slow decay will tend to consolidate commonly used tags and favour an emerging view of place, but it will be much more difficult for new tags to reach the new level of popularity as those already in the tag cloud. This is therefore a key parameter in setting the responsiveness of the system and the balance between a consolidated and a situated view of place.

The results from the interviews with users are presented in Table 3 organized by groups of questions.

Table 3. Survey summary (group results).

Questions group	Good	Sufficient	Poor	Bad
Relevance of the content presented in the display	48,8%	41,9%	7,0%	2,3%
Influence of the interactions in the display behaviour	30,2%	58,1%	9,3%	2,3%
Impact and acceptance	24,1%	58,6%	17,2%	0%
Handle with abusive behaviours	0%	83,3%	0%	17,2%
Interactions	32,5%	45,5%	15,6%	6,5%
Privacy	76,9%	0%	23,1%	0%

Overall, these results indicate that the perception of users in regard to the behaviour of the system was mainly positive. This is corroborated by the 90.7% of users that evaluate the relevance of the content presented in the display as Sufficient/Good.

The impact and acceptance of the system was evaluated as Sufficient/Good by 82.7% of users. The analysis of individual questions related to the “Impact and acceptance” shows that 92.9% of users consider that the content presented in the display is relevant to the context where the display is situated and 86.7% refer that the display contributes to enrich the place where it is situated, thus recognizing the utility of the system.

The various issues related to interaction with the system show that 78% of users evaluate as Good/Sufficient the mechanism of interaction with the display and all users acknowledged that the influence of their interactions is visible in the behaviour of the display.

In a specific question integrated in the set of questions related to Interactions, 84.6% of users refer that the system responded to their interactions in a timely manner. However, it was pointed out by some users (two cases) during the interviews that the system reaction was slow as they seemed to expect an almost immediate change in content when defining new tags.

In relation to “Handle with abusive behaviours” a significant portion of users (83.3%) stated that the display properly deal with abusive situations. However, 16.7% of users mentioned some offensive content presented on the display referring to the "small provocations." Although this has not been a central issue of research, and therefore, it is still necessary to carry out work to improve the behaviour of the display when subjected to offensive content requests, the proposed system has some features that allow lessen the occurrence of these situations. The presentation of offensive content is placed on two levels: the level of the Bluetooth device name and the level of content presented on the display. The Instant Places Service Module was a filter to the words used in the Bluetooth device name to block the interactions that contained "offensive words". The filtering is done using a list of “forbidden keywords” that is constantly updated. However, it is not easy to include all the words on this list and in some cases users specified words which although may be considered offensive they are not yet listed in that list. For the presentation of offensive content, which is originated in the interests specified by users, and although there have not been reports of these situations, there is the possibility of offensive content to be presented in the display. However, the filtering derived from the use of relevance indicators of the sources and the consequent exclusion of less significant sources can minimize the risk of using sources with more offensive content.

With respect to “Privacy”, 76.9% of users refer that the benefit from interaction with the display is higher than the cost of exposure of the Bluetooth device name. The unintentional exposure of the Bluetooth device name was prevented, since it is guaranteed to influence the content displayed on the display is necessary to use a specific command in the device name (*tag* in this case). So only the interactions of users that explicitly specify this command will be interpreted by the display as intentional interactions.

6 Conclusions

This paper has presented the results of a research programme on adaptive content for public displays. In our exploration of possible adaptation mechanisms, we started by approaching this issue as a context-aware problem. However, our first evaluations suggested that context *per se* was not enough to support adaptive content selection, in part due to the limited context information available, but also largely because of the challenges involved in specifying or inferring generic context adaptation rules for matching context situations with the utility of particular content items. More context information, and possibly more work on rule formulation, could probably lead to better results, but our subsequent work with the tag clouds would also highlight how a more simple approach, albeit one people can reason about, may also have important advantages. We have also identified obvious similarities between content adaptation in public displays and recommendation systems. However, there are also clearly different assumptions in both problems that severely limit the direct application of existing recommendations techniques.

Based on these observations, we have developed and evaluated a novel content adaptation system for public displays that uses a shared and public place profile in the form of a place tag cloud to combine the multiple interests of the place owner and place visitors. Overall, the positive results obtained during the evaluation suggest that this is a viable approach to the problem of content adaptation for public displays. The results showed that place visitors recognize the sensitivity of the system to their demands and that a place tag cloud was able to provide an adequate representation of place.

Despite their simplicity, tag clouds have revealed to be a promising solution to our specific scenario. Their main advantage is their ability to dynamically integrate the preferences of both place owner and place visitors into a single representation, and thus generate a representation of the social environment as a whole instead of each individual interest at a time. Moreover, they support a balanced combination of information filtering and information retrieval. They support information filtering because even if no one is using the system, the tag cloud is already there and capable to serve as a content generator. Also, the tag cloud specification defines an adaptation scope that limits the extent of the content that can be displayed. They support post-filtering because once the display is in place, people passing-by will sort out the content deemed more interesting to them. This way, the tag cloud will provide an interesting representation of the interests of a crowd. It avoids merely determining averages that are not representative and is able to deal with the tension between place adaptation, as something that can be learned over time, and situatedness, as the ability to react quickly to the social dynamics around the display.

Our intermediate studies have shown that simple keywords can be efficient content selectors, but are not always reliable representations of the concepts that people had in mind when proposing them. However, this can be considerably improved by framing user-generated keywords within the broader context of place, which may include location, organizational context, type of setting, or other keywords specifically created for that purpose. Even simple clues may be enough for overcoming the potential ambiguity in most keywords and guaranteeing an interpretation aligned with the nature of place.

The use of tag clouds as place models is obviously a limited approach from the perspective of descriptive richness and we are not suggesting that they are an appropriate model for other place representation scenarios. However, their simplicity and their visual nature may have been an important element in the positive results obtained with our system. Through the use of the tag clouds, we have made a key part of the system model explicit and visible, allowing people to reason about it and actuate according to their

goals. It also enabled people to perceive interactions being generated by other nearby users and observe the respective effect on the system. This ability to build an understanding of what was happening and interpret the system behaviour has been suggested in the interviews as something that may have contributed to increase user acceptance, even when the content shown was not perceived as the most appropriate option. We consider this to be an important finding, but it remains an interesting research topic to quantify this effect and be able to understand how does this ability to reason about what is happening in the subjective perception that people formulate about the system behaviour compares against other approaches potentially more effective and sophisticated, but less open to user inspection.

A final conclusion refers to the expectations that people may have about the nature of the system's reactions. We have always assumed the display to be mainly a communication medium through which people can express themselves in a way that becomes publicly available. Influencing the content presented on the display is one way for enabling the type of shared experience that public displays can afford. However, we perceived from interviews that peoples' expectations were not always aligned with this perspective. While the place-based tag cloud is essentially designed as crowd interaction mechanism, people often expect the system to exhibit an immediate reaction to their specific interaction. They seem to expect that an interactive public display should be there to provide them with the information they need, immediately and regardless of other people's interests. This notion of appropriating a public display for individual use, in a model similar to information kiosks or even web browsing with personal devices, is considerably different from the usage model we anticipated and may have generated some of the misunderstandings that emerged during the prototypes. This is understandable in a frame of reference in which interactive displays, even public ones, are mainly seen as devices for getting information and not so much as devices for publishing information. The observation that the system was slow to respond can be, to a certain extent, also interpreted as being a sign of this same frame of reference. In particular, an immediate response was never our goal. Still, a relatively slow responsiveness may represent a real limitation to the use of this approach in transient places.

Acknowledgments The first author was supported by a Portuguese Foundation for Science and Technology scholarship (SFRH/BD/31292/2006). This work was also funded by the Portuguese Foundation for Science and Technology (FCT) grant CMUPT/SE/0028/2008 (Web Security and Privacy).

References

- Arguello, J., Elsas, J.L., Callan, J., & Carbonell, J.G. (2008, March 30-April 2). *Document Representation and Query Expansion Models for Blog Recommendation*. Paper presented at the International Conference on Weblogs and Social Media, Seattle.
- Bihun, A., Goldman, J., Khesin, A., Marur, V., Morales, E., & Reynar, J. (2007). US Patent No. US Patent & Trademark Office: U. P. T. Office.
- Churchill, E.F., Nelson, L., Denoue, L., Helfman, J., & Murphy, P. (2006, August 1-4). *Sharing Multimedia Content with Interactive Public Displays: A Case Study*. Paper presented at the DIS2004, Cambridge, Massachusetts, USA.
- Das, A.S., Datar, M., Garg, A., & Rajaram, S. (2007). *Google news personalization: scalable online collaborative filtering*. Paper presented at the 16th international conference on World Wide Web, Banff, Alberta, Canada

- Dourish, P. (2004). What we talk about when we talk about context *Personal and Ubiquitous Computing*, 8(1), 19-30.
- González, G., Rosa, J.L.d.l., Dugdale, J., Pavard, B., Jed, M.E., Pallamin, N. (2006, 28 August- September 1). *Towards Ambient Recommender Systems: Results of New Cross-disciplinary Trends*. Paper presented at the Workshop on Recommender Systems in European Conference on Artificial Intelligence on Recommender Systems, Riva del Garda, Italy.
- Greenberg, S., & Roudind, M. (2001). *The Notification Collage: Posting Information to Public and Personal Displays*. Paper presented at the Human Factors in Computer Systems, Seattle, USA.
- Huang, E.M., Koster, A., & Borchers, J. (2009). *Overcoming Assumptions and Uncovering Practices: When Does the Public Really Look at Public Displays?* Paper presented at the 6th International Conference on Pervasive Computing Sydney, Australia.
- Interrelativity Inc. (2005, July, 22). Proactive Displays. Retrieved July, 2006, from <http://www.interrelativity.com>
- Izadi, S., Brignull, H., Rodden, T., Rogers, Y., & Underwood, M. (2003, November). *Dynamo: A public interactive surface supporting the cooperative sharing and exchange of media*. Paper presented at the Symposium on User Interface Software and Technology Vancouver, Canada.
- José, R., Otero, N., Izadi, S., & Harper, R. (2008, Oct.-Dec.). Instant Places: Using Bluetooth for Situated Interaction in Public Displays. *IEEE Pervasive Computing*, 7, 52-57.
- Leahu, L., Sengers, P., & Mateas, M. (2008, September 21-24). *Interactionist AI and the promise of ubicomp, or, how to put your box in the world without putting the world in your box*. Paper presented at the 10th international conference on Ubiquitous computing, Seoul, South Korea.
- McCarthy, J.F. (2000, April). *Active Environments: Sensing and Responding to Groups of People*. Paper presented at the Workshop Situated Interaction in Ubiquitous Computing at CHI2000, The Hague, Netherlands.
- McCarthy, J.F., Costa, T.J., & Liongosari, E.S. (2001). *UniCast, OutCast & GroupCast: Three Steps toward Ubiquitous Peripheral Displays*. Paper presented at the International Conference on Ubiquitous Computing, Atlanta, Georgia.
- McCarthy, J.F., McDonald, D.W., Soroczak, S., Nguyen, D.H., & Rashid, A.M. (2004, November 6-10). *Augmenting the Social Space of an Academic Conference*. Paper presented at the CSCW, Chicago, USA.
- McDonald, D.W., McCarthy, J.F., Soroczak, S., Nguyen, D.H., & Rashid, A.M. (2008). Proactive Displays: Supporting Awareness in Fluid Social Environments. *ACM Transactions on Computer-Human Interaction*, 14(4).
- McNaught, C., & Lam, P. (2010). Using Wordle as a supplementary research tool. The Qualitative Report. 15(3)(3), 630-643. <http://www.nova.edu/ssss/QR/QR15-3/mcnaught.pdf>
- Mehta, C. (2006). Tagline Generator. Retrieved September 2010, 2010
- Morales-Aranda, A.H., & Mayora-Ibarra, O. (2007, June 19 -25). *A Context Sensitive Public Display for Adaptive Multi-User Information Visualization* Paper presented at the Third International Conference on Autonomic and Autonomous Systems, Athens, Greece.
- Müller, J., Kruger, A., & Kuflik, T. (2007). *Maximizing the Utility of Situated Public Displays*. Paper presented at the Adjunct Proceedings of User Modeling, Corfu.

- Payne, T., David, E., Jennings, N.R., & Sharifi, M. (2006, August 29 - September 1). *Auction Mechanisms for Efficient Advertisement Selection on Public Displays*. Paper presented at the European Conference on Artificial Intelligence, Riva del Garda, Italy.
- Pessemier, T.D., Deryckere, T., & Martens, L. (2009). *Context aware recommendations for user-generated content on a social network site*. Paper presented at the 7th European interactive television conference, Leuven, Belgium.
- Ribeiro, F.R., & José, R. (2009a). Place-aware content selection from dynamic web sources for public displays *5th International Conference on Signal-Image Technology & Internet-Based Systems* (pp. 302-309). Marrakech: IEEE Computer Society Press.
- Ribeiro, F.R., & José, R. (2009b, March 23-26). *Timeliness for dynamic source selection in situated public displays* Paper presented at the 5th Int. Conference on Web Information Systems and Technologies, Lisbon.
- Russell, D.M. (2003). Large interactive public displays: Use patterns, support patterns, community patterns *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*: Kluwer Academic Publishers.
- Russell, D.M., & Sue, A. (2002, September 29 - October 1). *Using Large Public Interactive Displays for Collaboration*. Paper presented at the Workshop on Collaboration with Interactive Walls and Tables, Goteborg, Sweden.
- Schilit, B.N., LaMarca, A., Borriello, G., Griswold, W.G., McDonald, D., Lazowska, E. (2003, September 19). *Ubiquitous Location-Aware Computing and the "Place Lab" Initiative*. Paper presented at the WMASH'03, San Diego, California, USA.
- Seo, J., & Croft, W.B. (2007). *UMass at TREC 2007 Blog Distillation Task*. Paper presented at the Text Retrieval Conference, Gaithersburg, Maryland.
- Steinbock, D., Pea, R., & Reeves, B. (2007). *Wearable tag clouds: visualizations to facilitate new collaborations* Paper presented at the 8th International Conference on Computer supported collaborative learning, New Brunswick, New Jersey, USA
- Viegas, F., & Wattenberg, M. (2006). Communication-Minded Visualization: A Call to Action. *IBM Systems Journal*, 45(4), 801 - 812.
- Weiser, M. (1993). Some Computer Science Issues in Ubiquitous Computing. *Communications of the ACM*, 36(7), 75-84.
- Zhu, M., Zhang, D., Zhang, J., & Lim, B.Y. (2007). *Context-Aware Informative Display*. Paper presented at the International Conference on Multimedia and Expo, Beijing