

Place-aware content selection from dynamic web sources 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— Public digital displays could greatly benefit from the ability to dynamically select from the Internet content items that would be strongly related with the place where each display is installed. Generically, this is similar to the type of problem addressed by recommender systems. However, the usage context of a public display raises specific challenges that may limit the applicability of existing recommender systems. In this paper, we explore the creation of a recommender system for public situated displays that is able to autonomously select relevant content from Internet sources using keywords as input. This type of recommender system should enable public displays to become devices for Internet information delivery in public spaces, while also making them more situated in the social settings in which they are installed. We have created a recommender system based on these principles and we have conducted two studies to evaluate the perceived performance of the system. The results have shown that keywords can be very effective in driving user-generated content, but they often need to be complemented with contextual information that disambiguates their semantics.

Keywords- *Dynamic Sources, Information Integration, Public Displays, Situated Displays, Social Information Systems, Web Information Filtering and Retrieval, Web 2.0.*

I. INTRODUCTION

Public digital displays are increasingly common in all sorts of place. Despite this increasingly widespread use, studies have shown that most public displays are not very valued by their potential users [1]. They 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.

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. Still, the integration of that content in displays cannot be based on the same assumptions and principles that typically govern the definition of content in public displays. 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 that are evaluated as relevant for each particular display. 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.

Generically, the above problem corresponds to the type of problem addressed by recommender systems. Based on information about the preferences, past actions or properties of some entity, a recommender system is able to recommend new items that may be of interest to that entity. Similarly, a public displays could also take some type of place specification as the basis for dynamically selecting from web sources content that would be highly relevant for that place. However, the usage context of a public display raises specific challenges that may limit the applicability of existing recommender systems.

The first problem is generating some type of place profile that can be used as input to the recommender system. Rather than having a user profile, a public display needs a place profile that combines the preferences of the person managing the display, who we call the place owner, with the preferences of the multiple users that may use the display. This would lead to an evolving characterization of place that could drive the systems' behavior according to the social environment around the display. However, most public displays, even when supporting some form of interaction, do not generate meaningful traces of user activity. In most cases, the intended use is simply to be seen by the people around the display, meaning that they will normally be used without generating any information at all about how they are used, much less about the users interests or preferences. Without some type of information about user behavior we are not able to bring to public displays any of the approaches that have been so successfully applied on the web.

The second problem is the selection process after the recommender systems proposes a set of items. Most recommender systems assume some type of feedback or control by the user. The navigation experience gives people full control over which information to access and many cues on which information to select. Multiple data items from various sources are typically presented in the form of short summaries with links for further details, and people can easily evaluate which ones may be of interest and navigate accordingly. 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 skip to a next song if for some reason he or she does not want to listen to that song at that particular moment. In a public display, none of these possibilities is naturally available. People are very limited in their ability to influence the display decisions, not only for the technical considerations resulting from the lack of a mouse and keyboard, but essentially due to the fact that the display is public and shared. Furthermore, given that people will not normally have the possibility to request for further details, all the content is presented. This substantially raises the responsibility of the system in being able to determine the most relevant items from a pool of suggested content.

In this paper, we explore the creation of a recommender system for public situated displays that is able to autonomously select relevant content from Internet sources using a place specification as input. This type of recommender system should enable public displays to become devices for Internet information delivery in public spaces, while also making them more situated in the social settings in which they are installed.

This work is particularly focused on the viability of recommending content based on keywords specified by place owner. We are not specifically considering interaction scenarios, and thus in this case the place specification corresponds to a static weighted list of keywords defined by the place owner. However, in interactive contexts, this list of keywords could also be fed with contributions from users (as described in [2]) and thus the list and their weights would be dynamic. This is one of the important advantages that we associate with this approach even if in this study we are not specifically considering that scenario.

Our objective is to evaluate the viability of this keyword-based recommender process as a generator of relevant content for a situated display. We have created a recommender system based on the specific requirements of public displays, and we have conducted two studies to evaluate the perceived performance of the system. In a first study, done at the laboratory, we have asked participants to specify keywords and then evaluate the generated content according to the extent to which it matched what they expected when indicating a keyword. In a second study, in a real setting, we run a similar study, but this time the display was running for two weeks and at the end we collected a set of observations and logs.

The results have shown that keywords are simple to specify and are normally efficient as content selectors from dynamic web sources. They also have shown that an improvement can be achieved using contextual clues to give more semantic connotation to keywords.

II. RELATED WORK

This work builds on previous work in two main areas: adaptive scheduling in public displays and recommender systems in general.

Public display systems include some type of scheduling process that determines what gets presented when. 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. Content selection and space management on the display are typically resolved by users rather than by the system. BlueScreen [7] bases its scheduling decisions on history observation of exposed devices. These observations were collected during the advertising cycles of each ad in the past and also on the current set of detected devices which were in front of the screen. Muller [8] 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, current 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 based on users interests or user profiles. In this work we take a generic place specification and 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 are only addressing 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. Still, there are several specificities associated with this particular problem domain that are not common in most existing systems. Usually, approaches to recommender systems deal with applications that have 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 modeling 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 usually supplied by the source but in some cases this information is extended with other information that is inserted by users. These entities are the basis of all recommendation techniques. However, the way they are used can differ between different approaches. The three main recommendation techniques in use today are the content-based recommendations, the collaborative recommendations and the hybrid approach. In content-based recommendation, the system suggests to the user the items that best fit the user profile. The 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 the user preferred in the past. The adoption of this approach to our work presents some drawbacks that are commonly associated to this technique, 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. Despite some attempts to address news recommendation as a single collaborative filtering problem [9], particularities of public displays scenarios makes them unsuitable to our problem. These approaches make their 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. However, our target scenario is a single display and additionally, due the reduced number of users in our system and also due the inexistence of users ratings on past contents this approach is not suitable for this scenario.

III. A RECOMMENDER SYSTEM FOR PUBLIC DISPLAYS

Within the context of this study, we have developed a recommender system based on the specific requirements posed by content recommendation for public displays. This recommender system is based on a very simple place specification that basically corresponds to a set of keywords, possibly with differentiated weights which we call tag cloud.

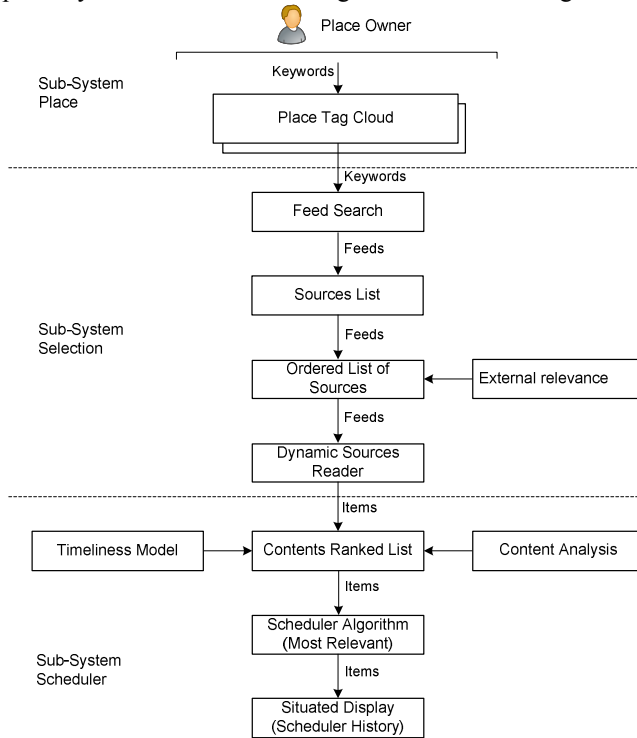


Figure 1. System architecture.

This recommender system takes the list of seed keywords, finds potentially relevant sources in the Internet, evaluates their immediate relevance according to multiple criteria and finally runs a dynamic scheduling process in which the most relevant items are cyclically presented. 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, Selection and Scheduler.

A. Sub-System Place

The Place sub-system represents the place preferences in the form of a weighted list of keywords, a tag cloud, corresponding to topics of general interest for the environment where the display is placed. Both, place owners and place visitors are beneficiaries of the display and thus they should be able to influence the display behavior according to their expectations, but in this work we are only considering the place owner role, who should specify the key guidelines for the display behavior. These include the keywords and a set of parameters associated with the scheduling process.

B. Sub-System Selection

This sub-system is responsible for the selection of a list of relevant sources using as base the place keywords.

1) Feed Search

Social applications represent an important resource to understand what users tend to be interested in. Using tags from the place tag cloud and respective popularity we search for related feeds on a news aggregator. Tags with higher values of popularity originate a higher number of sources.

2) Selecting Relevant Sources

Results from feed search are usually given by a large set of sources without any relevance criterion. After gathering contents from these sources a filtering algorithm is executed to avoid non relevant sources and to promote sources which content is well accepted by other users. Knowing what other people think about each source gives us an important marker that can help to improve the system utility preselecting the most relevant sources and eliminating sources that no one believes to be relevant. This is the reason why we define source self-relevance, a non contextual measure, i.e. it is independent of the usage context where the content is consumed. Important indicators for this measure can include the percentage of all Internet users who visit a given site or the traffic to the site.

Alexa search engine uses what they call a traffic rank. This measure is based on traffic data from Alexa Toolbar users. Moreover they provide an API that allows to access information about a specific Uri, namely their rank, their traffic and related sites. We use this rank that is country specific. Another important measure of feed relevance is the number of users subscribed to the feed. Higher values of subscriptions denote higher feed interest. Through Newsgator API we get for each relevant source the number of users that subscribe it.

Because both traffic rank and feed subscriptions have distinct numeric domains they need to be normalized. To do this we distinguish two intervals: one where the display manager considers the quantity as good; the other where the display manager considers the measure as bad. Sigmoid functions characterize such intervals and provide a smooth interpolation between the limits of those intervals. Fig. 2 and Fig. 3 illustrate the equations for the traffic rank relevance function (TR) and for the subscribers' relevance function (SR).

$$TR(Rank) = \frac{1}{1 + e^{\frac{good+bad-2Rank}{good-bad}}}$$




Figure 2. Traffic Rank relevance (TR).

$$SR(Sub) = 1 - \frac{1}{1 + e^{\frac{good+bad-2Sub}{good-bad}}}$$




Figure 3. Subscribers Relevance (SR).

Both traffic rank and number of subscribers are combined in a function that defines the external and non contextual self-relevance of each source.

Finally, the dynamic sources reader is responsible for read from the web the selected list of sources.

C. Sub-System Scheduler

Evaluating the relevance of a particular source in the context of a public display could involve complex inferences about current situation. In this work, we specifically use three criteria: keyword weight, timeliness and content.

1) Keyword Weight

This criterion considers the relative weight in the place model of the keyword that has been used as the search key when selecting the proposed content.

2) Timeliness

This criterion is concerned with the time dimension, i.e. evaluating how timely the information is. Timeliness represents the temporal pertinence of the content and changes over time even if the respective source remains unchanged. This notion of timeliness is of an obvious importance in setting the relevance for any type of source. However, 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. In a separate piece of work [10] we analyzed a set of sources from distinct types (e.g. news, events, announcements, blogs, etc.). Based on that analysis, we propose two timeliness formulas for two common types of source, those based on a publication date and those based on a planned event date.

3) Content Analysis

Some heuristics are used to analyze the relevance of the body content. Because content from selected sources may include distinct media content like: text, images, video, etc., and it may be provided with distinct properties in their fields (e.g. length, type of contents), they should be analyzed before being displayed to the users. This analysis is based in three heuristics: text analysis, link analysis and image analysis. Text analysis includes format conversion and content description analysis. Information about content like

the description length and content language is then used to influence the content relevance. Content with reduced description is penalized in their relevance (e.g. description lengths lower than 50 characters has only 20% of relevance). Link analysis, is performed to analyze the number of links in the content description is also used to influence the relevance of the content. For example a content description that contains many links is penalized in its relevance. This penalization is defined considering the relation between the number of characters that belongs to Url and the total description length. Image analysis, is performed with the major purpose of extracting image links. This information is further parsed allowing images being adapted for presentation (e.g., adapt the image size).

4) Next Content Selection

While in many recommender systems it is supposed for a content to be presented once, in public displays there are some content that is supposed to be presented more than once, depending on its relevance and the target population. Additionally, the scheduler should adapt its decisions to the dynamic and heterogeneous social environment around the display and it needs to balance between the place owner control and the interests of place visitors. Thus, content selection is realized in three steps. First, source relevance acts as a threshold level for accepting or dismissing each source. Sources with lower values of external relevance are dismissed and are not considered as content providers. Second, place model selects the next topic that represents the most relevant topic at the moment. Third, from the set of contents that belongs to the same cluster of the selected topic, the relevance of each content item is calculated as a function of timeliness and content relevance, and considering the current defined scheduler behavior (e.g. do not repeat the same item two times consecutively).

Finally the most relevant one is selected to be presented.

This algorithm is cyclic and it is constantly verifying the place for the most relevant keywords. A keyword that is currently relevant may be non relevant in the next period, and later, it may be relevant once more. Thus, a content that is presented at the moment may be at a later date newly relevant and thus presented once again.

IV. EXPERIMENTAL WORK

To evaluate this type of recommender system we have conducted two separate experiments. The first one evaluates the appropriateness and accuracy of the method for obtaining relevant sources using as base the set of seeds. The second one is a complimentary experiment to evaluate how methods studied in the first experiment are used and perceived by users in a real scenario.

The evaluation goals are organized around four main groups: method simplicity, method appropriateness and accuracy, some insight into typical patterns for seed specification and relations between place owner perception on source relevance and other external relevance indicators. For each group the following evaluation goals were defined:

- 1) *Simplicity*:
 - (S1) Place owners are able to successfully specify its preferences using simple keywords.
 - (S2) Place owners consider that their preferences can be described using a small set of keywords.
- 2) *Appropriateness and accuracy*:
 - (A1) Place owners consider that the system responds appropriately to their specifications.
 - (A2) Place owners consider that sources provided by the system are adequate to their seeds specifications.
- 3) *Patterns on seed specification*:
 - (P1) There are insight typical patterns for seed specification on different place owners.
- 4) *Relations between place owners relevance perception and external relevance*:
 - (R1) Place owners' perception on source relevance and external source relevance.

B. *Experiment I - Obtaining Relevant Sources Using as Base a Set of Seeds*

This experiment was realized to evaluate the appropriateness and accuracy of the method for obtaining relevant sources using as base the set of seeds. Participants were interviewed and asked to take the role of place owner for specifying their preferences for sources to be displayed in a situated display.

1) *Experiment Description*

This experiment involved 25 participants (researchers and students in different computer science areas) and consisted in the following steps:

a) We asked participants to play the role of place owners and we explained to them what the purpose of the experiment was and what was expected from a place owner in the definition of the behavior of the public display. We then asked them to specify their preferences using at maximum four keywords.

b) Based on the keywords specified, 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. These feeds are the same for all the participants regardless of the respective keywords and they serve as a ground truth for our study.

c) Each participant seated in front of a computer display where it will be shown content from the feeds. For each feed three items were shown, but without the feed title and without any indication of the originating keyword. Each participant was shown 30 sources. 20 of them are derived from the place owner specifications and equally distributed among all the specified keywords, except when a keyword did not select any source. The other 10 sources are randomly selected from the predefined list of sources. Participants were told that there were multiple selection algorithms being used, and that they could expect some of the content to be totally unrelated. However, 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. Therefore, participants never knew what was the origin of each feed being presented.

d) For each feed presented, participants were asked to classify to what extent the respective content corresponded to their expectations. They 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. The answers were given directly in the system through a button in the same window where content is being displayed. Every time a new source is presented and evaluated a set of information is stored. This information includes: source title, source Uri (both stored but not shown in the display) source title and description of three items of the source, traffic rank, number of subscribers, keyword that originates the source, a string that indicates if the sources is from a seed or from the predefined list and the user evaluation. The system presents the next source only after the user gives the answers to the currently displayed content.

e) To evaluate the appropriateness of the keywords as a means for expressing place owner interests and preferences, we run a simple questionnaire at the end of the content presentation part, with the following four questions:

- How simple is the preferences specification using simple keywords?
- How appropriate is preferences specification using simple keywords?
- What is the number of keywords that you consider appropriate for preferences specification?
- Are there other methods that you consider appropriate for preferences specifications?

The two first questions are multiple choice questions. Possible answers for the first question are: very simple, simple, hard, very hard. For the second question possible answers are: not appropriate, slightly appropriate, appropriate and very appropriate. Question three is answered using intervals: [0,3], [4,6] and 7 or more. Finally, the fourth question is an open and optional question. Additionally we use information from questionnaire jointly with information collected during the experiment to study three other questions:

- Some insight into typical patterns for seed specification (e.g. similar keywords).
- System ability to generate, based on initial seeds, content that matches the expectation of the person who specified those seeds.
- Relations between place owner perception on sources relevance and external relevance measures.

2) *Results and Analysis*

Twenty five participants completed the experiment. A summary of the collected data is presented in table I.

TABLE I. SUMMARY OF COLLECTED DATA.

Users	25
Keywords per user (Average)	3,3
Keywords (Total)	83
Keywords (Distinct)	55
Sources presented (Total/From keywords)	726/493
Sources presented (Distinct)	327
Suggestions on open question	7
1 st Keyword most referred	Informatics-8 times
2 nd Keyword most referred	Sport-7 times

80% of all place owners specified their preferences easily and they obtained appropriate or very appropriate results in more than 67% of presented contents. However, some participants (about 4%) seemed to have failed to understand correctly how to specify their preferences through keywords and they specify a set of keywords with no clear meaning and in some cases using verbs (e.g. go). 16% of place owners indicated in the open question that it is important that the system should allow the definition of preferences using more than one keyword or using operators between keywords. For example, this is the case of an user that wanted to specify “Hello Kitty” and other that wanted to specify “action games”. Despite referred difficulties more than 72% of obtained results are related or very related to their specifications. Additionally, answers on question 3 of the questionnaire shows that 96% of users refer that their preferences can be specified appropriately using at maximum 6 keywords. This value validates premise (S2): Place owners consider that their preferences can be described using a small set of keywords.

With respect to the appropriateness and accuracy of the method for preferences specification, table II presents a summary of the evaluations.

TABLE II. SUMMARY OF THE EVALUATION.

	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%)

There are 726 evaluations on source appropriateness. 233 of them correspond to evaluations of sources deriving from the predefined list and 493 correspond to sources obtained through users' specifications. 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.

These results show that place owners are able to successfully specify their preferences using simple keywords (premise S1) but the use of more than one keyword per preference can improve results.

Additionally, question 1 in the questionnaire shows two important things. First, methods for interests specification using keywords are considered appropriate by users to describe their preferences (88% evaluate them as appropriate or very appropriate), i.e. place owners consider that the system responds appropriately to their specifications (premise A1). Second, seeds preferences specification allows

to obtain appropriate sources (63,3% are considered as related to their specifications) what validates the premise (A2) “place owners consider that sources provided by the system are adequate to their seeds specifications”.

Table III presents a description of 50% of all keywords specified by users. People tend to be interested in topics that are related to their job/occupation and their hobbies. We make the experiment in a school of technology with computer science courses and results show that 21,7% of all keywords are related to the school and its teaching areas and 48% of users specify at least one keyword related to the place where they work/study. The second most referred keyword is “Sports” (7 times/ 8,4%). However, if we include all related keywords in a main topic “sports” it represents 16,9% of total specified keywords (14 times).

TABLE III. FREQUENCY OF KEYWORDS SPECIFIED BY USERS.

Specified 8 times	Informatics
Specified 7 times	Sports
Specified 3 times	Photography, Cars
Specified 2 times	Technology, Education, Benfica, News, Economy
Specified 1 time and related to Technology	Internet, PDA, Computers, Programming, Apple, Mobil-phone
Specified 1 time and related to Sports	Football, Squash, Ball, Paintball, Tennis
Specified 1 time and related to News	Actuality

To evaluate the premise (P1) “There are insight typical patterns for seed specification on different place owners” a set of guidelines should be considered: Users tend to indicate topics that are related to their job/occupation; General interests like sports or news are also indicated; Hobbies are also an important topic that is many times referred to. These guidelines suggest that, although we are not able to identify the main hobbies of users, we are able to identify a set of interests that corresponds to the preferences of an important part of them. In fact, a description of preferences that includes topics related to the place where the display is situated and some general interests like sports and news can represent the preferences of many users (64% in the case of this experiment). This emphasizes the situated nature of these displays. However, remaining preferences are very difficult to identify because they are user specific and they present much diversity.




Figure 4. Evaluation on appropriateness vs. external relevance indicators.

Fig. 4 shows users evaluations on appropriateness and the relevance value derived from the traffic rank and from the number of subscribers. The graph shows that there is a relation between users' evaluations and external indicators

on relevance. Sources that are evaluated as *related* or *very related* present higher external relevance values than sources that are evaluated as *nothing appropriate* or *no opinion*. This confirms that these indicators can be used to refine the list of sources to use (Premise R1).

C. Experiment II – Evaluate the System in a Real Scenario

To support some evaluation outside the controlled environment of the lab, we also conducted a second experiment in which our main purpose was to evaluate the system in a more realistic scenario.

1) Experiment Description

To specify place interests we invite four people of the Information Systems Department (DSI) general office to play the role of the place owner.

The experiment consists on the following tasks:

a) We asked to each place owner to specify his interests in the form of keyword seeds. Each one is asked to specify a set of ten keywords that represents topics they think to be interesting to be presented in the display.

b) We gave to each place owner a diary on which to register any relevant events or simple stories about the system behavior.

c) The system generates content based on the place specification obtained from all contributions, during a period of two weeks.

d) A final interview was made to assess user perception on the display behavior and the overall appropriateness of presented content according to their expectations.

2) Place Specification

Each place owner specifies a set of ten keywords which represent content topics he wants to see in the situated display. We use four people as place owner for the same situated display to obtain a richer description of the place interests and also to find relations between their specifications. A superficial analysis of the set of keywords specified by the four place owners shows that the majority of them are related to the place itself and activities that takes place on it. Other set of keywords is related to the place, but many of them only makes sense if they are contextualized or associated to other words (e.g. park, place, etc.). Finally, there is other set of keywords that is related to hobbies or personal interests (e.g. holidays, music, etc.).

3) Setting

This experiment takes place in the hall of the general office of the DSI. This is a place where people wait for one's turn on the general office and it is also a transient space where people go through to their office or lab. Users are mainly teachers, staff, students and researchers in technological areas. There are some periods in which the number of people passing by is higher than normal (early in the morning; and before and after lunch). During the remaining period the number of people near the system is relatively small (maybe a few dozens per hour).

The situated display has two presentation areas: one at the top of the display that presents the set of keywords that represents the place interests. The popularity of each keyword is represented by the keyword size. The second presents the content that is selected by the scheduler. At the top of this pane a queue of keywords shows keywords that originate previous two contents, the present content and the next two contents that will be displayed.

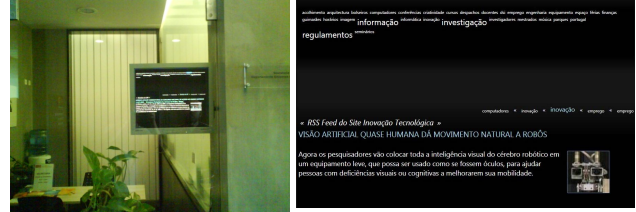


Figure 5. Left: display in the real scenario; Right: display layout.

There are two distinct periods on the experiment. Each period corresponds to seven days of the experiment. The difference between these two periods is related to the use of distinct scheduler rules. These rules influence both the source relevance and the behavior of the display. Table IV presents the set of rules used in each period. The first two rules influence the display behavior. Higher values of these rules make less frequent the presentation of the same item and less frequent the presentation of items from the same source and thus, sources and items that never were presented in the first experiment period will be presented in the second experiment period. As a consequence content with lower values of relevance that does never appear in the first period can be selected to appear in the second period. The last rule influences both the source relevance and the display behavior. In the second period the display presents a more dynamic behavior because content is selected from a broader range of topics and sources.

TABLE IV. SCHEDULER PARAMETERS.

Behavioral rules	1 st period	2 nd period
Do not repeat the same source within a set of	1	10
Do not repeat the same item within a set of	5	15
Number of items to consider per source	2	1

4) Evaluation

Place owners filled in a diary with their opinions. The register diary is composed by an initial introduction with a small explanation about what is intended to do and then a reserved space for each day of the experiment. Place users are informed in advance about how to fill the register diary. The task of filling up the register is spontaneous and without any constraint. Each place owner is free to write his judgments about the situated and about contents and its appropriateness.

5) Results and Analysis

Results are presented in two phases. First we present some statistical data about content presented in the two parts of the experiment. Secondly, we analyze data collected from users' evaluations obtained during the experiment period and also after the experiment period through the interview. Table V presents data about the system organized in each one of

two periods. Data is restricted to the period 8am to 8pm, and weekend days are excluded.

TABLE V. STATISTICAL DATA ABOUT CONTENT IN THE DISPLAY.

	1 st Period	2 nd Period
Number of schedules	9899	10040
Number of distinct items presented	259	279
Number of distinct sources presented	29	44
Number of keywords that originates items presented in the display	9 (25,7%)	14 (40,0%)
Average value of the relevance of items presented in the display	65,0	21,8
Average value of the timeliness relevance	70,8	47,6
Average value of the content relevance	84,2%	70,8%

Data from the register diary of the four place owners was completed with a final interview to better understand their perception on contents appropriateness. 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 (even though the respective feeds may be very popular). Several other anecdotal examples were similar in that they also corresponded to clear mismatches between the initial intention associated with the keyword and the generated content. Some keywords when taken in isolation may be interpreted in many ways, which seems to point out at the need to contextualize some keywords giving them a more concrete semantic connotation.

V. CONCLUSIONS

Keywords are simple to specify but their efficiency as seeds for user-generated content may vary greatly. Rather than leading users into more complex conceptual formulations, for example by asking for keyword pairs or by proposing ontologies, we think we may considerably improve the selection process by framing user keywords in a broader place context defined by physical location, organizational context, type of setting, or even the overall set of user-generated keywords. In previous work [11] we argued that these contextual elements per se were not enough to support adaptive content selection because it was far from obvious what type of association rules could be created between them and particular content. That was part of the motivation for exploring keywords as an alternative and more intentional path for context-aware displays. This work has confirmed that keywords are normally efficient as content selectors, but it has also shown that they are not always reliable as representatives of the concepts that people had in mind when proposing them. As a result, we conclude that the combination of the two may offer the most promising approach. Even simple contextual clues may be enough in disambiguating ambiguous keywords and in providing an interpretation aligned with the nature of place.

As future work we expect to extend our recommender system with this ability to frame keyword-based selection within the context provided by the additional contextual information associated with place.

ACKNOWLEDGMENT

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

REFERENCES

- [1] 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.
- [2] 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.
- [3] J. F. McCarthy, D. W. McDonald, S. Soroczak, D. H. Nguyen, and A. M. Rashid, "Augmenting the Social Space of an Academic Conference," presented at CSCW, Chicago, USA, 2004.
- [4] 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.
- [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] 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.
- [10] 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.
- [11] F. R. Ribeiro and R. José, "Proactive Scheduling for Situated Displays," in *Workshop on Ambient Intelligence Technologies and Applications*. Guimarães, 2007.