City Research Online

City St George's, University of London

Citation: Goker, A. S., Cumming, H. & Myrhaug, H. I. (2004). Content retrieval and mobile users: An outdoor investigation of an ambient travel guide. Lecture Notes in Computer Science, 4200, pp. 23-31. ISSN 0302-9743

This is the unspecified version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/615/

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Content Retrieval and Mobile Users: An Outdoor Investigation of an Ambient Travel Guide

Ayse Goker
School of Computing
Robert Gordon University
St. Andrews Street
Aberdeen, UK
asg@comp.rgu.ac.uk

Hannah Cumming
School of Computing (SmartWeb),
Robert Gordon University
St. Andrews Street
Aberdeen, UK
hc@comp.rgu.ac.uk

Hans I. Myrhaug
SINTEF Information and
Communication Technology,
N-7465 Trondheim,
Norway
hansim@sintef.no

ABSTRACT

People's information needs change as they encounter new situations. The need for an ambient information environment becomes more evident in the case of the mobile traveller where situated information access is one of the main challenges.

The motivation for this work has been to provide relevant information to the right situation and user in an ambient manner. Our way to solve this is to deliver personalised and context-aware information to the mobile user. To this end we have developed a platform, and prototype applications for travellers, and tourists. The system integrates our own tag technology with information from content service providers covering both general travel guide and local information.

The development methodology is user-centred, iterative, and progressive in nature. It combines information retrieval (IR) test and evaluation techniques with iterative and user-centred development techniques at the test and evaluation phase. Combining the two disciplines gives us the ability to test and evaluate both the *information* aspects and the *interaction* aspects of any information system in parallel. Another advantage would be that one can develop content and software in parallel.

This paper focuses on the IR test and evaluation framework that has been used in conjunction with the user-centred development. We emphasize the importance of performing IR test and evaluation for mobile systems in terms of users' situations and tasks. The paper presents the results of some of the findings from a preliminary user test in an outdoor scenario. The test took place in a popular tourist destination in Spain.

Keywords

User-oriented test and evaluation, ubiquitous content access, mobile information retrieval.

1. INTRODUCTION

The trend in miniaturization continues as with the growing amount of handheld and embedded computers. Until recently, few applications and content services were developed and made available on such equipment. However, this is already changing as we can observe the trend amongst telecom operators trying to channel content on top of providing network access. Digital content is, therefore, fully on its way into handheld and wireless networked devices. This seems to lead to an increased need for ubiquitous information access. From the user's perspective, the thrust is about having access to and receiving relevant information to the situation. Seen from an information retrieval perspective, it is about populating people's pockets and the surroundings with intelligent search engines that operate on ubiquitous and distributed content repositories. One implication of this might be that digital information will be rendered close to the body on handhelds, personal belongings, and in clothes - or embedded inside objects, furniture, rooms, and open areas in the surroundings – ready for its distribution.

Our approach in addressing the challenge of information needs is by enhancing the environment with: vicinity-based content

Figure 1: The tag developed and a mobile device with the content

service access and delivery, search engines for mobile devices, and context-aware technology. We demonstrate how one can provide people with information experiences, whilst being mobile, by distributing information relevant to the local situations. The aim is for relevant content to automatically be distributed and delivered onto mobile phones and other handhelds in the vicinity of physical objects, rooms, and open areas (see Figure 1).

The focus in this paper is on presenting the results of the first user test iteration of an ambient travel guide.

2. BACKGROUND AND RELATED WORK

Related work can also be found in the fields of ubiquitous and context-aware computing. Dey et al, in a special issue on Situated Interaction and Context-aware computing [8], provide an overview. The focus from this perspective, however, has tended to be on location-based approaches and device contexts. Examples of these can also be found in few applications for tourists. Wider perspective of context has been discussed in some forums e.g. [13, 28]. Tourist applications have been developed by several for research purposes [1, 6, 9].

Within IR research, contextual information provides an important basis for identifying and understanding users' information needs. Cool and Spink in a special issue on Context in Information Retrieval [7] provide an overview of the different levels in which context for information retrieval interest exists. They refer to information environment level (e.g. Taylor [26]), information seeking level (e.g. Belkin [2]), information retrieval interaction level (this refers to user-system interactions but from a cognitive perspective can be said to relate to Ingwersen's cognitive communication model [15]), and the query level (e.g. as also discussed within TREC lately).

These categories are related and overlap. To this extent, the work described here has aspects in each of the four categories, but the first three in particular. The query level parts are not based on a linguistic analysis of the query but a case of augmenting or expanding it with contextual information. Others have viewed context-aware retrieval more as a way of filtering results from normal retrieval techniques [4].

More specifically, however, related previous work involved the development of a context learner for a probabilistic information retrieval system [11] in a traditional environment with a bibliographic search system. This was based on observations (within that environment) that users will tend to repeat searches or conduct a series of closely related searches over a period. Whilst each search must be regarded as representing a different information need they could be said to occur within a particular context.

At present, the more general user context model [19] has a wider range of aspects which have been enriched as a result of the earlier work in context-aware computing [18]. This work is based on the overall perspective of providing the right information, at the right time, to the right user i.e. providing information to the right situation or context. Thus, as a concept, the user context model represents aspects of the user's situation spanning from social, physical, and spatio-temporal aspects to user interests.

Much work in ubiquitous computing and ambient intelligence has focused on pervasive applications e.g. smart house, airports with varying size and functionality of display units, and WIFI. Considerable work in context-awareness focuses on the location-based applications [5, 14] devices, and networks [10]. Some work is done on the interaction between the user and device [25]. However, little is done on the nature of the information that users might receive in ambient digital environments and how to interact with systems delivering the information.

3. OVERALL SYSTEM DESCRIPTION

We have developed an ambient and context-aware information system. The overall system includes three cornerstones: a specifically developed *tag* (context tag), *a content service provider*, and the *user* with a *mobile device* – Figure 2. The system integrates our own tag technology with information from content service providers. The content is from both a general travel guide publisher and local information providers..

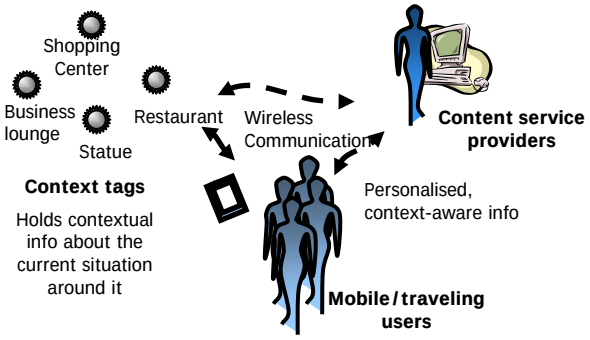

Figure 2: Overview of system architecture

The diagram illustrates the flexible ways in which users can access information. Content service providers may provide online information directly to a user (usually at a significant cost to the mobile user) or also via tags mounted in various strategic places thus creating an information zone. Information can be uploaded from remote via WLAN or Ethernet (by content service provider or building owners, for example) and be accessible locally to the user who is in that environment and situation. For example, in the context of an ambient travel guide, the historic and cultural web pages, local sights, shops, maps, and local events can be communicated to the mobile phones. Web-pages and other multi-media content can be relayed and distributed via the tags. Further details of the system and reference architecture developed and implemented can be found in Myrhaug et al. [20]. Demo of some of the applications has been presented by Goker et al. [12].

The system enables a user to perform the following information seeking tasks:

- Browse the content via categories
- Search for content via queries to a search engine
- Walk around in the surroundings and automatically receive content pop-ups within particular information zones

The user can also receive recommendations or personalized search results based on preferences. These may be explicitly

stated, or implicitly derived through assessing their search behavior.

The system supports the simultaneous execution of several of the above tasks at the same time. E.g. the user can walk and receive content from the surroundings whilst also browsing or searching. A usage scenario is given below in order to illustrate the possible ways a user can seek and be notified of information in such systems.

Example Usage Scenario: The user is at a hotel reception. The local tourist information office has mounted tags at different sites in the city, including some hotels. When the user is in the vicinity of a tag, he has the option of receiving a small software client which enables communication with the tags. He accepts the installation, and the thin client is installed. The client can receive, present, and cache up to 1Mb of content, whilst the total amount of content available for the digital travel city guide is 20Mb. 60% of it consists of images, and the rest are (pointers to) pages optimized for mobile phones.

The user now leaves the hotel and starts to discover the city. A map is always available on the phone so he can find other places to cache relevant content. He performs a search and receives some 10 results, of which 5 of them are available on the phone and the remaining ones can be received if he encounters the vicinity of a tag. Two minutes later he stands in front of a cathedral where a tag is mounted in the surroundings. The missing 5 pages followed by a pop-up item specifically on the cathedral are automatically transferred to the mobile device.

He is not really interested in cathedrals, so he specifies to the mobile client that he is interested in events and shopping. The system automatically deletes some of the cached content on the mobile, and pre-caches more relevant content onto the mobile. He starts to walk while intermittently browsing some interesting events. The tag and the mobile phone disconnects as he moves further away. A bit later on, he encounters a third tag, which is mounted on a shopping centre. He automatically receives a pop-up of special offers in the shopping centre as he stands there.

4. EXPERIMENTAL SETUP AND METHOD

Experiments and evaluation effort within mobile and ubiquitous environment has largely has had gravity around the mobile devices and the user interfaces. Generally, the information items themselves have been of less focus. We have aimed to address this point.

We address the evaluation of the information items through *situated experiments*. With situated experiment we mean an experimented where the user interaction is located (i.e. situated) both physically and mentally in daily/ normal situations of system use. Hence, it is related to the concept of situated interaction referred to in human-computer interaction and ubiquitous computing.

Early work in conducting situated experiments in IR can be found in work referred to in [11]. Task-oriented views have been considered important [21, 27] and other notions of relevance judgments addressing the situation such as situational relevance have also been discussed (described in [3]). It is now increasingly accepted to plan for situated experiments.

Our test and evaluation approach goes further by combining these with user judgements in the real situation - or as close to the 'real' situation as possible. Hence, one aim is to provide an IR test and evaluation framework that can be integrated with traditional user-centred development methods – and in particular made for mobile and ubiquitous information systems.

We have, therefore, conducted several situated and task-driven user experiments, both *indoor* and *outdoor*, to help evaluate our mobile information systems, and make progress on our IR test and evaluation framework. Our method is being developed so that it can be used for experiments during both real and controlled system use within mobile IR.

We generally believe that experiments in Mobile IR should be as *strongly situated* as possible. However, we also realise that sometimes in order to control various variables, a more weakly situated experiment might be necessary. Situations can be *real* or *controlled* —and naturally there are degrees of control. For example, we conducted the experiment described in this paper in Seville, capital of Andalusia region in Spain, at locations where many tourists visit. It was conducted during times of the day when travel guide information is quite likely to be sought. Thus, this particular experiment is *more strongly situated* than conducting it in a remote research lab, and asking the users to imagine they are at the location. Location, time of day, and user groups, are example variables that can affect how strongly situated an experiment is. Another important parameter is the user's task. Tasks can be real, as performed by users going about their daily activities, or they can be specified to the users. Naturally, the descriptions can be closer to the likely needs of the user group or somewhat remote from their usual information needs. The latter would involve more cognitive effort and imagination on the part of the user.

In this experiment, we controlled the *tasks* but aimed for these to be as realistic as possible for the users' situations. The paper presents some of the results from a preliminary test of an outdoor scenario in a popular tourist destination. The reader can observe that user-centred test and evaluation techniques were used together and in parallel with the IR oriented test and evaluation techniques.

The experiment had several purposes including the testing of our technical platform. However, the main purpose of the experiment was to:

- Help assess the mobile information retrieval system by noting the extent to which users are satisfied with the content items they are shown in certain situations
- Gain an insight into users' information needs before, and while they travel
- Obtain general feedback on the system but more importantly on the information seeking tasks mentioned in section 3 (browsing, searching, pop-up while walking)

4.1 Experimental Environment and Overall Design

The experiment was conducted in Seville - a tourist destination ideally suited to many kinds of tests and evaluation of mobile and ubiquitous systems for tourists and travellers. The actual test area was the touristy city centre with a part also coinciding with some traditional businesses in the centre.

The test area starts at the gates of a famous castle (Alcázar) and stretches to small square (Plaza del Salvador). It has been chosen because it contains many diverse points of interests such as: shops, hotels, monuments, bars, restaurants, meeting places, attractions, historical buildings etc., catering to a large variety of interests users might have. Additionally, all this is in a relatively small and dense space, which has been tailored for pedestrians and tourists – an aspect which sets this area apart from other possible ones, making it more suitable for user-oriented evaluation.

Tags were mounted at a variety of places in the centre and a walk-about route involving passing through several *tag zones* (described shortly) was planned. The route was designed to link up a series of plausible information needs that might occur during a walk around the city centre. This was done in conjunction with relevant tourist expertise. Each tag location had at least one main attraction or potentially interesting places (points of interest near it).

We refer to the area defined by the signal range of the context tag as an *information zone*. Users entering an information zone can have access to certain kinds of content. It is also possible, for some items to pop-up on the user's mobile device when encountering such zones.

4.1.1 Situated and Task-oriented Information Retrieval Approach

We focused this experiment on 6 locations: 5 separate information zones (defined by location of tags) and one warm-up location for brief system interaction training. In each information zone, we conducted two tasks (see below). There was also one introductory warm-up task for each user. There were 5 users, thus each performing 11 tasks – a total 55 tasks. In addition to the tasks, users were given pre- and post-questionnaires, together with informal interviews.

The tasks were split into two types: specific (S) and general (G). The S tasks were to test a specific feature such as the pop-up of a content item when a user is within an information zone. The G tasks gave the user a more open task, such as to find information on a topic, site, restaurant, and so on but something plausible given the situation. These searches for information were performed via a query box on the mobile device. Although the user could browse, this task was not tested in this preliminary test. Much of information retrieval work on the mobile device stresses the use of categories to traverse through on the grounds that it is easier for the user to click through a series of short menus. However, here we wanted to investigate the use of our search engine in a mobile environment. We did not test our category-based view in this particular experiment.

Our *search engine* for the mobile device produced a ranked ordered list of content items, such as those one sees in web search engines. For each of the top five content item rank positions, the user was asked to give *relevance* (topical relevance) and *usefulness* judgments (in relation to the user's current situation), along with where in the ranked list the user thought the content item should have appeared.

The tasks were given during a walk around the city centre. The locations were 1) a hotel (very centrally located) where system was introduced; 2) a restaurant (highly recommended for its food and atmosphere); 3) a shop (watch shop near a popular square); 4) a religious shop (selling numerous statues and dresses for religious festivals – in particular an upcoming one); 5) a cloth/textile shop (potentially en route between two squares); 6) an antique shop (neighbouring a shopping centre). Locations no. 2-6 had tags, while location 1 had no tag in the surroundings in order not to reveal the Specific tasks prior to walking. However, general searching was possible at location 1.

As described earlier, each of the information zones had two kinds of tasks associated to it: A specific task (S1-S6) set towards assessing a content item pop-up; and a general task (G1-G6) involving wider information search exercise. Each location ties to a number between 1 and 6. Hence, with G1 we refer to the General task performed at location 1. With S6, we mean the Specific tasks performed at location 6. The S tasks result in assessing only one particular item (the content pop-up in situation) but the G tasks involve assessing the top 5 items in a ranked ordered list, after the user has entered a query. As the G tasks, help evaluate the search engine, we focus more these tasks in the paper.

Below is the list of general tasks (G1-G6) given to users; the locations they were given during the walk-about city centre; and an outline description of the task set.

Task Description Summary Extracts

The purpose of G1-G6 was to test the system effectiveness with respect content searching task (2nd bullet in section 3). In short, the tasks communicated to users were:

G1: … you are hungry, want to eat and are looking for food/ a restaurant (in the area of Calle Alemanes)… need information…

G2: …need information on or about the Cathedral.

G3: … need information on or about Plaza de San Francisco.

G4: … need information on or about Plaza del Salvador.

G5: …would like to find interesting places around Calle Francos.

G6: … want to find some information about the Peyre Shopping centre, and shops within it.

4.1.2 Pre and Post Test Questionnaires

A pre and post questionnaire was also given to the users before and after the tasks, described above, respectively. The Pre Questionnaire was used to gather information about the user, their computing experience and their travel habits. The Post Questionnaire was used to gather overall views on the usefulness of the system and the content after a user test had been completed. These were also followed by informal interviews in order to give feedback on the system design and the content. Hence, the user testing followed a traditional user-centred and iterative development approach.

4.2 Measures

This work has been influenced by the approaches to relevance as described in Saracevic's earlier work [22-24]. We have used this kind of approach in the Web domain but have extended it for the mobile and ubiquitous search environments. Mizzaro [17] gives a comprehensive account of the discussions and approaches to relevance while also stressing the importance of situational relevance and the task that the user is trying to achieve.

We obtained relevance judgements as Relevant (R), Non-Relevant (N), and Partially Relevant (P). The work done at University of Tampere [16] involving graded relevance judgements was considered but not used. We stress the importance of obtaining relevance judgements at the situation the user is in and not retrospectively or whilst in another situation. In mobile information access this is particularly challenging as there are several considerations to bear in mind. The user is a traveller/ tourist and time is of the essence. Asking for 'relevant' or 'not relevant' and even asking about the 'usefulness' seemed an acceptable question load per task. Stretching this to incorporate 'partially relevant' was in fact at times a stretch and was more clearly asked when the user seemed hesitant about saying a content item was relevant or non-relevant. It seemed that the graded scale approach would simply be too complicated to ascertain under the conditions we wanted our experiment to work in.

4.3 Content Collection

Figure 3 below shows a version of the system designed for Sony-Ericsson P900. The figure shows a search statement being entered in a search box, and the corresponding list of retrieved content items. The user can subsequently click on any and view a full description of each item.

The underlying content collection comprises both general travel guide and local content for Seville. The general travel guide was an appropriate extract provided by Lonely Planet Publications. The local content was provided by Sevilla Global, which is owned by the City Council of Seville.

Both data sets are in XML, and they were the basis for creating the indexes.

4.4 Experiment Procedure and Technical Setup

The purpose of the system is to integrate and demonstrate our tags with information from content service providers (and subsequently also deliver personalised, context-sensitive information wirelessly to the handheld device). However, only a "plain" IR search engine was tested as part of this preliminary investigation. Essentially this was based on a standard TF*IDF (term frequency * inverse document frequency) approach in information retrieval. Any user context information was not added at this point.

Figure 3: Screenshots of content on mobile device. Ayuntamiento (city council) is the main tourist attraction immediately adjacent to Plaza de San Francisco

The tag developed is a miniaturised computer, more than an RFID tag, It can hold 128 MB of content and is Bluetooth enabled. It can also be WLAN enabled and be updated remotely via Ethernet. However, in this experiment we used the tags as stand-alone computers not connected to Ethernet or WLAN.

The mobile devices we have used in our iterative experiments are: HP iPAQ and Sony Ericsson P900. The content on the tags can be pushed to or pulled from the mobile device and presented on the screen - it can be delivered remotely from a content service provider through the wireless network infrastructure, or via the tag.

The users started off on their test trail from a hotel lobby where they were introduced to a test person. The test person gave some introductory information and the preliminary questionnaire. The users were told that there would be a walk-about where they would be given some tasks at specific locations. Prior to starting, however, they were shown and briefly trained on the main features of the mobile device. During the walk-about, they were guided by a test person who described the tasks as they arrived at the specific locations. The test person also recorded their judgements of the content items shown (based on the user's feedback). This was done in order to ensure the user had both hands free to use the mobile device, and also to avoid interfering with the task flow and walk-about.

5. EXPERIMENTAL RESULTS AND ANALYSIS

The results are described in two parts: results related to the information retrieval tasks; and results of the questionnaires followed through with informal interviews.

5.1 Results for the Situated and Task-oriented Information Retrieval Tasks

Let us, for example, consider the user relevance judgements for the top two rank positions for General task G1 and Users 1–5 (U1–U5). These were:

U1: R, R; U2: R, R; U3: N, N; U4: R,N; U5:N, -.

In other words, U1 found both of the top 2 content items Relevant as did U2. U3 on the contrary thought none of the top 2 shown was relevant. U4 thought the 1st was but not the 2nd. U5 did not find the 1st relevant and did not answer the rest explicitly (as saw no point after the first judgement in that case). Hence, the precision at the 1st rank position is 60% (3/5). In other words, 60% of content items at this 1st position in the ordered list were considered relevant by the users. The precision when we consider the top 2 content items (i.e. cut-off at rank position 2) was 62.5% (5/8 as we did not count the last judgement for U5. There were no Partially-Relevant judgements in this sample extract.

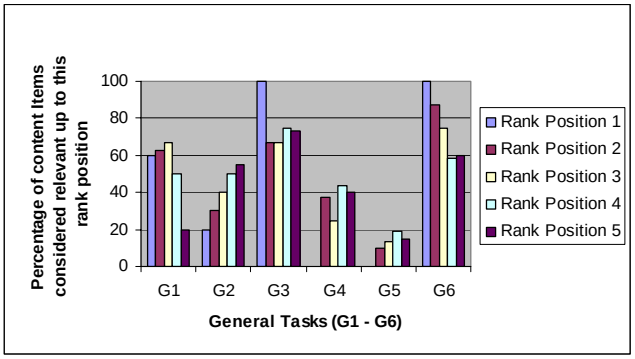

Figure 4: Precision Values for the General Tasks with only Relevant Content Items

Figure 4 shows the percentage of the content items considered relevant so far, for each rank position. Some results were quite poor with G5 receiving < 20% considered relevant at any of the ranked positions and the 1st item returned was not considered relevant by any of the users. G5 was a general task (info around Calle Francos) and results were for shops at a number of the top ranked positions. Most users commented that they were not looking for shop info at that time and this explains why the task received low scores. Results for G4 were also low for similar reasons. G2 involved a general query about the cathedral and some of the top results were restaurants close to the cathedral. Users found the restaurant information irrelevant. In general, they considered restaurant content items to be irrelevant for the specified task but would have considered them relevant had they been hungry.

G3 and G6 achieved best relevance results with each of the position 1 ranked items considered relevant by all the users and at all other rank positions about 60% were considered relevant. G3 (Plaza de San Francisco) resulted in a high number of historical type content items which the users preferred and this is reflected in the high relevance results. Interestingly, although G6 involved finding information on a shopping centre (Peyre), the scores were high – at odds with results of the other tasks where shop type information was received. However, the informal interviews after the field test showed that the users are *not totally adverse to shop type information* but rather that the timing is more critical –if explicitly in shopping mode than the system received higher scores.

Results improve when content items that were considered partially relevant were included as relevant, an increase from 46% to 68% of items. 67% of all items were also considered to be useful. The comments the users provided during the tasks gave more information on how they made their judgments and what type of information they want. One common comment through out the tasks was that the information was relevant to their location but not to them, to explain why they only gave the content item a partial relevant rating and not a relevant rating.

G1 was the warm up task to find a restaurant close the hotel. This task received high relevance results at the first few ranking positions but the relevance falls off quite quickly at later ranking positions. This is simple because the first few positions were restaurants as the users expected while the later results were about their location but not relevant to their query. All the evidence stress that the context of the user is very important when returning

results, the subjects did not like information that was not relevant to them even if it was relevant to their location.

5.2 Results of the Pre and Post Questionnaires and informal interviews

Before starting the tasks users were asked questions about their travel habits. This was used to help provide background information for interpreting some of their feedback during the information search tasks they were given. Subsequently, post questionnaires were also used to help highlight what their general attitude towards this way of interacting with the system was. This was followed by informal interviews.

The subjects were in their twenties making them potentially familiar with mobile technology. This, however, varied in practice. All used mobile phones daily and rarely used PDAs. Familiarity with wireless technology was mixed: one familiar with all and two not knowing any. Figure 5 shows what type of travel information was gathered before and while travelling. As expected sites/attractions, events, maps, are always important and is the type of information needed both before and during travel. Other travel information such as accommodation and transport are needed at different times –before and during travel, respectively.

One interesting result is that shop information was not sought before travelling and even more unexpectedly not explicitly sought while travelling either – important considerations when deploying the tags.

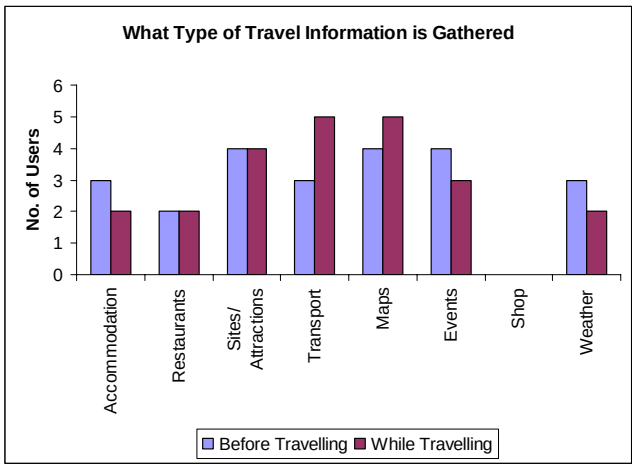

Figure 5: What type of travel information users gather before and while travelling

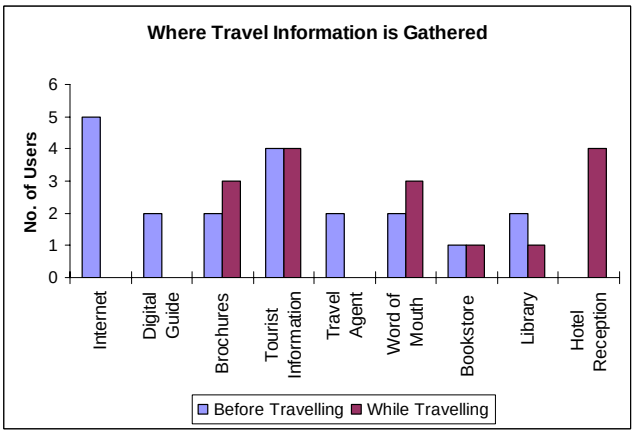

Figure 6: Where travel information is gathered before and while travelling, sum of all users

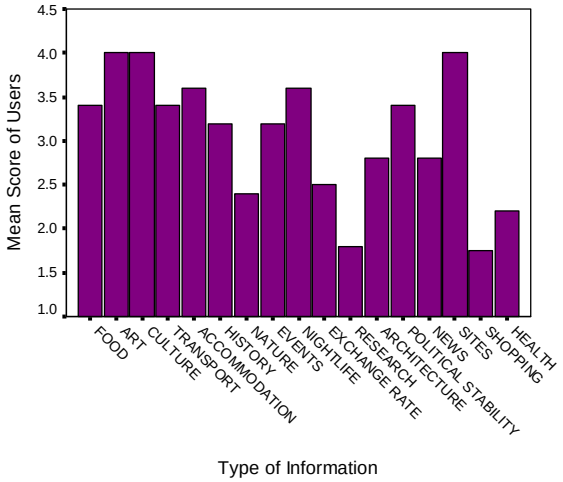

Figure 7: The importance of each type of information averaged for all test users

For the above types of information, the internet is the main source of travel information before travel, while during travel there is a complete change and hotel reception becomes the primary source (see Figure 6).

Users were also asked to expand on the previous questions to find out how important each type of travel information is to the subjects. The subjects were given a five point scale with which to rate the different travel information, where 1 is unimportant and 5 is important. Figure 7 shows the mean score for each type of information across all subjects. These results support those from the previous question showing that the highest rated type of information was sites, art and culture. Also considered important information was transport, accommodation, nightlife, food and political stability. The types of travel information that were considered unimportant were research, health, nature, exchange rate and shopping. Although the category types are not exactly the same, the results coincide with the previous question with sites and events scoring high and shopping scoring low.

6. Discussion

The comments the users provided during the experiment gave more information on how they made their judgments and what type of information they wanted. One common comment throughout the tests was that the information was relevant to their location but not to their situation (this was a plain version of the system with no user profile incorporated or system adaptivity), explaining why they only gave the content items, a partially relevant rating and not a relevant rating.

Shop information was not wanted during general queries where users were looking to find information about sites. One user commented that they would like a way to specify whether or not to include shop type information. One user commented that instead of an overall description about the shop, he would have preferred specific product details from the shops. Restaurant type information was generally considered more relevant than shop type information with a couple of users stating that they would have preferred restaurant items rather than shops items. Time was also considered important, the restaurant information was considered more relevant at lunch time. This again shows the importance of the user context on how they make their relevance judgments.

Other observations the users made were on the amount and type of information the content items contained. As was seen in earlier results the users were generally happy with the amount of information provided by the content items. However, the amount of information considered enough was different for the different types of content items. Additionally, the scores for amount and structure of the information were lower in the general tasks than the specific tasks due to the lack of information on some points of interests. Some items only had short descriptions where as the content for the specific tasks was generally more complete.
Hence, considering the user's situation (context) is important both when testing and when designing the content and the interactive system components for mobile use.

6.1 Conclusion and Further Work

In post task questionnaires and follow up interviews, users repeated that they would prefer more historical/site type information to more shops. However, there were some potential indications as to when shop information might be suitable but further experiments are needed before concluding on this. They also wanted maps to help find the places the system recommended –rather unsurprisingly.

Results showed that the users knew what to do and they checked the PDA regularly for pop-up web pages. This was not the only method of delivery; users were able to search when they wanted too as per the searches they performed during their tasks. However, it does seem that they were positive and comfortable with this means of interaction with the surroundings for mobile and ubiquitous information access. When asked to give a score rating on how the system would benefit the user personally, the mean result was 4.6, 5 being the maximum benefit. All users in this small sample felt that they were better prepared for exploring the city using the system.

Further work is being carried out in the city with a larger number of users, a wider range of tasks, and improved search engine algorithms.

7. ACKNOWLEDGMENTS

This work has been funded under the EU-IST FP5 project AmbieSense (2001-34244). Our thanks to Murat Yakici and Ralf Bierig (RGU) who helped conduct majority of experiments. We are grateful for the content provision and local organisational support of Johan Vängbo (Sevilla Global), content provision and expertise of Ron Gallagher (Lonely Planet Publications). Marius Mikalsen (SINTEF) and Stein Tomassen (SINTEF) for user interface programming on earlier versions of mobile device and Sree Kanth Nuti (RGU) for later versions. Thanks also to Stuart Watt and Chris McKillop for their comments and feedback on earlier versions of this paper.

8. REFERENCES

[1] Abowd, G..D., Atkeson C.G., Hong, G., Long S., Cooper R., and Pinkerton M. Cyberguide: a mobile context-aware tour guide. Wireless Networks. 3:5. Special issue: Mobile computing and networking: selected papers from MobiCom '96., 421-433, 1997.

[2] Belkin, N.J., Anomalous States of Knowledge as a basis for information retrieval. Canadian Journal of Information Science 5, 133-143, 1980.

[3] Borlund, P. The concept of relevance in IR. Journal of the American Society for Information Science and Technology, 54:10, 913-925, 2003.

[4] Brown, P.J., and Jones, G. Context-Aware Retrieval: Exploring a new environment for information retrieval and information filtering. Personal Ubiquitous Computing 5(4): 253-263, 2001.

[5] Chen, G. and D. Kotz. A Survey of Context-Aware Mobile Computing Research. Dartmouth, Dept of Computer Science, Dartmouth College: 16, 2000.

[6] Cheverst K., Davies N., Mitchell K., Friday A., and Efstratiou, C. Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences. In Proceedings of CHI 2000, Netherlands, April 2000, pp 17-24.

[7] Cool, C., and Spink:, A. Special issue on context in information retrieval. Information Processing and Management 5, 38 (2002).

[8] Dey, A., Kortuem, G., Morse, D., and Schmidt, A., Special Issue on Situated Interaction and Context-Aware Computing, Journal of Personal and Ubiquitous Computing, 5, 1, (2001).

[9] Dunlop, M., Morrison, A., McCallum S., Ptaskinski P., Risbey, C., and Stewart F. Focussed Palmtop Information Access Combining Starfield Displays with Profile-Based Recommendations. Mobile and Ubiquitous Information Access, Lecture Notes in Computer Science series 2954, 79-89, Springer-Verlag, 2003.

[10] Ferscha, A. and Mattern F. Pervasive Computing, 2nd International Conf Pervasive 2004, Proceedings, (Eds), Linz-Vienna, Austria, Springer-Verlag, 2004.

[11] Goker, A. Context learning in Okapi. Journal of Documentation (53), 80-83, 1997.

[12] Göker, A., Yakici M., Bierig, R., and Myrhuag, H.I. A Context-sensitive Information System for Mobile Users (Demo). Workshop on Information Retrieval in Context, SIGIR 2004, Sheffield, UK, 2004. http://ir.dcs.gla.ac.uk/context/

[13] Gross,T., and Specht, M.Workshop on User Modeling for Context-aware Applications. User Modeling Conference UM2001, Sonthofen, 2001.

[14] Harter, A., A. Hopper, et al. (2002). "The Anatomy of a Context-Aware Application." Wireless Networks 8: 187-197.

[15] Ingwersen, P. Information Retrieval Interaction. London: Taylor Graham, 1992 (0-947568-54-9).

[16] Kekalainen, J., and Jarvelin, K. Using graded relevance assessments in IR. Journal of the ASIST, 53:13,1120-1129, 2002.

[17] Mizzaro, S. Relevance: The Whole History. Journal of the American Society for Information Science, 48:9, 810-832.

[18] Myrhaug, H.I. Towards Life-long Context Spaces. In: Workshop on User Modeling for Context-aware Applications. User Modeling Conference UM2001, Sonthofen, 2001.

[19] Myrhaug, H. I., and Goker, A. AmbieSense - interactive information channels in the surroundings of the mobile user. 10th International Conference on Human-Computer Interaction (HCI International 2003)}, Crete, Greece, (2003).

[20] Myrhaug, H., Whitehead, N., Goker, A., Faegri, T.E., and Lech, T.,C. AmbieSense – A System and Reference Architecture for Personalised Context-Sensitive Information Services for Mobile Users. 2nd European Symposium on Ambient Intelligence EUSAI 2004, Eindhoven, Netherlands, 2004. Springer-Verlag (to be published).

[21] Reid, J. A task-oriented non-interactive evaluation methodology for information retrieval systems. Journal of Information Retrieval, 2:1, 115-129, 2000.

[22] Saracevic, T. Relevance: A review of the literature and a framework for thinking on the notion in information science. Journal of the American Society for Information Science, 26, 321-343, 1976.

[23] Saracevic T., and Kantor, P. A study of information seeking and retrieving. I. Users, questions, and effectiveness. Journal of the American Society for Information Science, 39, 177-196, 1988.

[24] Saracevic T., and Kantor, P. A study of information seeking and retrieving. II. Users, questions, and effectiveness. Journal of the American Society for Information Science, 39, 197-216, 1988.

[25] Situated Interaction in Ubiquitous Computing, Workshop, CHI 2000, Netherlands.

[26] Taylor, R.S., Information use environments. In 'Progress in communication sciences'. Eds Dervin B., Voight, M. 217-255. Ablex: Norwood, NJ., 1991.

[27] Vakkari, P. A Theory of the Task-based Information Retrieval Process. Journal of Documentation 57 (1): 44-60, 2001.

[28] Workshop on Information Retrieval in Context, SIGIR 2004, Sheffield, UK, 2004. http://ir.dcs.gla.ac.uk/context/