Who, What, Where & When: A New Approach to Mobile Search

Karen Church
Adaptive Information Cluster
University College Dublin
Belfield, Dublin 4
karen.church@ucd.ie

Barry Smyth
Adaptive Information Cluster
University College Dublin
Belfield, Dublin 4
barry.smyth@ucd.ie

ABSTRACT
Mobile devices and the mobile Internet represent an extremely challenging search environment. Limited screen-space, restricted text-input and interactivity, and impatient users all conspire to exacerbate the shortcomings of modern Web search. Recently researchers have proposed that typically vague search queries be augmented by context information, as a way to help search engines to retrieve more relevant information. In this paper we propose a novel interface to support multi-dimensional, context-sensitive mobile search, combining context features such as location, time, and community preferences to offer a unique search experience that is well-adapted to the needs of mobile users.

ACM Classification Keywords
H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval, H.5.2 [Information Interfaces and Presentation]: User Interfaces

General Terms
Design, Experimentation, Human Factors

Author Keywords
Mobile Search, Search Interfaces, Mobile Web, Context, Preferences, Personalisation, Location, Time, Communities

INTRODUCTION

According to many commentators, the world of mobile information access has recently turned a corner as important new developments in devices, services, and billing models are set to usher in a new age of anytime information access. Certainly, the introduction of fully fledged mobile information access devices, such as Apple's iPhone and Nokia's N95, represent an important step forward in terms of functionality and usability. In addition, the major Internet players such as Google, Yahoo, MySpace and Microsoft, to name but a few, have recently set their sights on mobile users with the introduction of a new generation of mobile-ready services and applications. And finally, mobile operators are facilitating a more active mobile Internet subscriber-base by introducing a new range of flat-rate billing models that serve to encourage rather than penalize high-volume usage.

Against this backdrop, search engines will continue to provide the basic information access infrastructure that subscribers will need to locate relevant information in a timely manner. Indeed there has been considerable recent activity in this regard with all of the major search engine players launching the latest mobile editions of their Web services. In the main, however, mobile search offers little more than traditional Web search, with most players simply adapting their Web interfaces to meet the display characteristics of mobile handsets. Recent research has highlighted the shortcomings of this traditional approach to search in a mobile context, with increasing evidence that mobile searchers are frequently failing to locate what they are looking for in a timely manner; vague queries continue to dominate mobile search and few users have the patience to search long-result lists on their mobile handsets [3, 4].

Over the years many researchers have argued the need for search engines to be more responsive to the preferences of searchers and the context of their queries. In our work, we believe that providing context-sensitive, personalized search is critical to the success of mobile search but this view has only been adopted in a very limited form by current search engines. For example, Google's new mobile search[1] and Yahoo's oneSearch[2] both offer users with access to location-sensitive search, ensuring that the searcher's physical location is factored into the selection and ranking of results.

No doubt these improvements with help users to search more effectively but we believe that they represent the tip of the ice-berg when it comes to what needs to be done. In particular, we argue that these improvements are ultimately limited by the traditional query-based search interface on which they are based, one in which context and preferences play an ad-hoc role in guiding search, and in this paper we argue the need for a more radical re-imagining of search on the mobile Internet. Specifically, we propose that context and pref-

[1] http://arstechnica.com/news.ars/post/20070327-hittin-the-streets-with-googles-new-mobile-search.html
[2] http://mobile.yahoo.com/mobileweb/onesearch

IUI'08, January 13-16, 2008, Maspalomas, Gran Canaria, Spain.

erences must play a more dominant role in mobile search and we present a new *search browser* interface. In this new interface we will integrate context, in the form of temporal and location information cues, with preference information derived from the queries of like-minded communities of mobile searchers. In addition, we will break with the traditional *query-box-result-list* view of search, presenting the searcher instead with a view of evolving search activities that is sensitive to their context and preferences. In this way, prospective searchers can *browse* through community search experiences and manipulate the searches of others, learn from these searches, and initiate their own, in a way that takes full advantages of the strengths of their mobile handsets.

RELATED WORK

Traditional Mobile Search Interfaces

One area of research related to the focus of this paper is the very specific challenge of how best to present search results on a mobile device. Presentation and interface design is critical in mobile services and traditional Web search approaches can be problematic in the mobile space given the obvious high screen real-estate demands that such approaches impose on mobile devices. In some of our earlier work we examined an alternative approach to presenting search results on a mobile device, called *result-gisting* [2, 4]. Result-gisting enjoys the informativeness of snippet text while providing for a more economic use of limited screen real-estate. The core idea behind the approach is to replace result snippets with a much shorter text representation that is made up of the terms of related queries that have led to the selection of a particular result in the past. The approach was made possible as a direct consequence of a community-based personalised search engine called I-SPY, which records the queries and result selections of communities of users [9]. Results from our previous evaluations of the result-gisting technique indicate that related queries do provide an informative alternative to snippet text and they also suggest that users judge the use of related queries to provide a better balance between informativeness and screen-space on mobile handsets [2, 4].

In [7], Jones et al. use a set of key phrases instead of standard snippet text. The idea behind this approach is that these key phrases will provide informative result surrogates in the face of uninformative result titles. These key phrases are automatically extracted from the result pages. Their evaluations show that the resulting key phrases provide for a more economic use of screen space and are at least as effective and informative as using long result titles.

Alternative Mobile Search Interfaces

Another area of research directly related to this paper is alternative approaches to mobile search interfaces that steer clear of traditional Web search techniques. FaThumb [8] is a user interface designed for navigating through large data sets on mobile devices providing a more efficient means of mobile search. FaThumb uses faceted metadata navigation and selection as well as incremental text entry to narrow the results. The metadata facets are displayed in a 3x3 grid. Each facet contains a label and a number representing how many records can be characterised by that label. A user evaluation indicated that the traditional text-entry approach is faster if the query was something specific, like a name for example, however, if only the data characteristics are known, facet based navigation is faster.

The Questions not Answers (QnA) system [6] also provides an interesting alternative to traditional search interfaces. Rather than examining how to provide high-quality search results, i.e. the *answers* to queries, their approach is to provide access to previous queries, i.e. the *questions*, relating to a user's location. This novel interface displays queries made by other people in a given location using a map-based interface, providing users with an enriched sense of place. By clicking on the queries users can execute the displayed search. In a recent evaluation of the QnA prototype [5], users found the interface to be very useful and they enjoyed the increased level of interaction it enabled.

Our current work is similar in nature to the QnA approach. The QnA system essentially *tags* queries with a location. These queries are displayed on a map-based interface enabling users to visualise the search space. The QnA prototype does not, however, provide any means for a user to filter queries, other than by location. Given that the volume of queries at specific locations is likely to be quite high and there is no means to filter queries, the QnA prototype raises a new interface/presentation challenge. We think this is a core area to address given the unique characteristics of the mobile space.

Therefore, our new mobile search interface utilises contextual information such as location and time as well as preference information, derived from the queries of like-minded communities of mobile users, to provide a unique search experience. Unlike traditional query-based search, which requires user input before providing any information to end-users, our new *search browser* interface provides mobile users with information immediately thus encouraging *discovery* of new content. The interface presents historical query and result-selection data for users to browse through on an interactive map-based interface. Users can filter queries and result-selections based on context and preference information. In the following section we will describe the key features of our prototype in more detail.

PROTOTYPE

We developed a simple mobile search engine prototype to illustrate our proof-of-concept search interface. The prototype uses a set of queries and result-selections from two different sample communities, one representing *entertainment* and the other community representing *tourist sites*, all within the central Dublin area.

How it works

When the user first initialises the application, the user is shown a map centered at their current location (Figure 1). The map shows all *recent queries* and *result-selections* submitted by other users in that specific location. Queries are shown by a red marker while result-selections are distin-

guished by a yellow marker. We refer to these queries and result-selections as the *prime set*. The map is updated periodically so that newly entered queries and result-selections are also displayed.

Figure 1 illustrates the initial search interface shown to the user. The interface shows a small number of queries and result selections scattered at various locations on the map. In this case the queries are a mix of entertainment related queries like *Events Pravda, 35 Lower Liffey Street Dublin 1 Ireland* and more tourist-type queries such as *Dublin castle Great Ship Street Dublin 2*.

Figure 1. Our new mobile search interfaces showing queries and result selections made by other users in a given location.

Figure 2. When the user manipulates the time slider, the display is updated accordingly.

At the bottom of the interface there are two sliders. Users can interact with these sliders to filter queries and result selections. The first slider represents *time*. This slider ranges from *now* to *earlier*. The default value for the time slider is *now*. When the default value is selected the interface plots all queries executed most recently. Users can adjust the time slider to essentially go back in time and display queries and result selections submitted at an earlier date. These queries might have been submitted yesterday, last week, last month, etc. The further back in time the user chooses, the more queries and result-selections displayed on the interface. An interesting feature of the time slider is that it allows users to visualize the temporal patterns of queries within a given location. Users can easily visualise query histories and how queries have changed/evolved over time. Figure 2 shows the updated interface after the user manipulates the time slider. We can see that as the slider has been moved towards the *earlier* marker, more past queries are displayed.

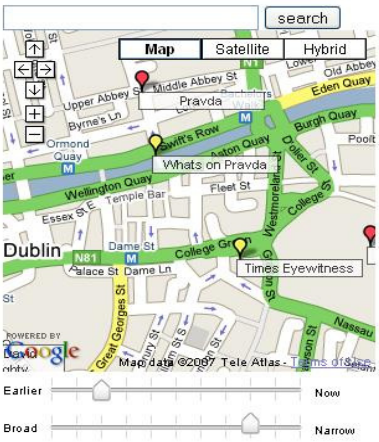

Figure 3. When the user manipulates the preference slider, the display is updated accordingly.

The second slider represents user *preferences*. This preference data is derived from the past queries of like-minded communities of users. The preference slider is labeled *broad* to *narrow* which refers to *broadly focused on the community* vs *narrowly focused on the community*. The default value for the preferences slider is *broad*. Again users can manipulate the preference slider to *adjust* the set of *prime queries* shown on the interface. We use a form of *query similarity* to model the preference context. When a user accesses the application, the system automatically calculates the *similarity* between the user queries and all other queries in the dataset. Using a simple measure of query overlap (also known as query similarity) from [1], we can measure various degrees of overlap between queries, ranging from duplication of terms to partial overlaps.

Manipulating the preference slider modifies the similarity threshold and thus relaxes or tightens the community aspect of the interface. Figure 3 shows the updated interface after the user manipulates the preference slider to display queries *narrowly focused on the community*. We can see that as the slider has been moved towards the *narrow* marker, less queries and result-selections are displayed. In this case the user in question is interested in *entertainment* type queries. That is, the user has previously executed entertainment related searches and is thus part of the *entertainment* community. Therefore as the preference slider is moved towards narrow, less similar queries are omitted from the interface.

As well as allowing users to filter queries based on key contextual and preference data, the interface also allows users

to interact with individual queries and result-selections. As mentioned previously, queries are illustrated by a red marker with an associated label while result-selections are illustrated by a yellow marker with an associated label. Clicking on the *query marker* causes the displayed search to be re-executed (See Figure 5). Clicking on the *result-selection marker* opens a small information window. This window displays the most popular/recent results selections, i.e. result URL's, as well as a link to execute the query in question. Users can choose to go directly to one of the listed URLs or they can choose to re-execute the query (See Figure 5).

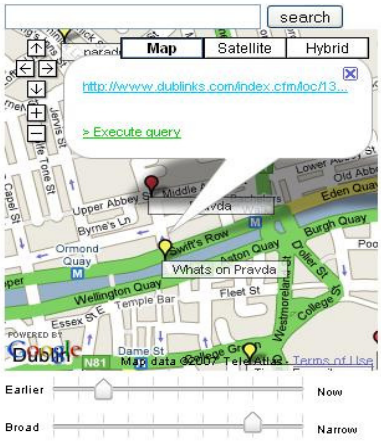

Figure 4. Clicking on the result-selection marker opens an information window showing the most popular result selections.

Query: ***Whats on Pravda***

PravdaVision: Wallace & Gromit
PravdaVision has become a popular Monday fixture in Dublin's city centre. Classic films, free popcorn and a relaxed atmosphere have lured most Dublin film fanatics! The first film is Wallace & Gromit ...

Figure 5. Results of executing the search *Whats on Pravda*

CONCLUSIONS & FUTURE WORK

In this paper we have presented a novel proof-of-concept interface to support multi-dimensional, context-sensitive mobile search, combining context features such as location, time, and community preferences to offer a unique search experience that is well-adapted to the needs of mobile users. Unlike traditional search interfaces, which require user input before providing information to end-users, our interface gives mobile users interesting information from the beginning. The interface presents historical query and result-selection data for users to navigate through on an interactive map-based interface. The rich UI enables users to interact with the interface, execute searchers, view past result-selections and filter queries based on context and preference information.

We are currently investigating a number of different areas relating to our new mobile search prototype. Firstly, we are in the process of implementing a fully-functional working prototype. We plan to deploy the prototype on an iPhone or similar mobile device and evaluate the system with real mobile users in various locations around Dublin. We are also exploring the use of communities in this application and how to best model the social side to mobile search. At present we are using a query similarity threshold derived from the past queries of like-minded users, however, there are other interesting approaches we'd like to investigate further. Finally, we plan to investigate the use of result-selections as place descriptors in more detail. The use of result-selections is one of the key novel elements of our interface and therefore this is an important area for us to explore as part of future work.

ACKNOWLEDGEMENTS

This material is based on works supported by the Science Foundation Ireland under Grant No. 03/IN.3/I361.

REFERENCES

1. E. Balfe and B. Smyth. An Analysis of Query Similarity in Collaborative Web Search. In *ECIR '05: Proceedings of the 27th European Conference on Information Retrieval*, pages 330–344, 2005.
2. K. Church, M. T. Keane, and B. Smyth. Towards more intelligent mobile search. In *IJCAI-05: Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 1675–1676, 2005.
3. K. Church, B. Smyth, P. Cotter, and K. Bradley. Mobile information access: A study of emerging search behavior on the mobile internet. *ACM Transactions on the Web*, 1(1):4, 2007.
4. K. Church, B. Smyth, and M. T. Keane. Evaluating interfaces for intelligent mobile search. In *W4A: Proceedings of the 2006 international cross-disciplinary workshop on Web accessibility (W4A)*, pages 69–78. ACM Press, 2006.
5. M. Jones, G. Buchanan, and D. Arter. Incidental information and mobile search. In *MobileHCI '07: Proceedings 9th International Conference on Mobile Devices and Services*, pages 155–158. ACM Press, 2007.
6. M. Jones, G. Buchanan, R. Harper, and P.-L. Xech. Questions not answers: a novel mobile search technique. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 155–158. ACM Press, 2007.
7. S. Jones, M. Jones, and S. Deo. Using keyphrases as search result surrogates on small screen devices. *Personal and Ubiquitous Computing*, 8(1):55–68, 2004.
8. A. K. Karlson, G. G. Robertson, D. C. Robbins, M. P. Czerwinski, and G. R. Smith. Fathumb: a facet-based interface for mobile search. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 711–720. ACM Press, 2006.
9. B. Smyth, E. Balfe, J. Freyne, P. Briggs, M. Coyle, and O. Boydell. Exploiting query repetition and regularity in an adaptive community-based web search engine. *User Modelling and User-Adapted. Interaction*, 14(5):383–423, 2004.