

Towards More Intelligent Mobile Search

Karen Church, Mark T. Keane and Barry Smyth

Adaptive Information Cluster,
Smart Media Institute, Department of Computer Science,
University College Dublin, Belfield, Dublin 4, Ireland
{Karen.Church, Mark.Keane, Barry.Smyth}@ucd.ie

Abstract

As the mobile Internet continues to grow there is an increasing need to provide users with effective search facilities. In this paper we argue that the standard Web search approach of providing snippet text alongside each result is not appropriate given the interface limitations of mobile devices. Instead we evaluate an alternative approach involving the use of related queries in place of snippet text for *result gisting*.

1 Introduction

The mobile world is growing at an astonishing rate. In 2004, the number of mobile users worldwide reached 1.52 billion, in Europe alone the number of users exceeded 340 million and the volume of pages viewed on the Mobile Web continues to rise [CellularOnline, 2004]. Five years has passed since the publication of the renowned WAP usability report, which detailed the various limitations of both WAP (as a means of mobile Web access) and mobile devices [Ramsay and Nielsen, 2000]. However, the Mobile Internet has since been successfully re-branded by leading mobile operators like Vodafone and O2, thus encouraging users to take advantage of next-generation mobile devices.

If we examine current mobile devices, we can divide them into three classes. At the lowest level we have standard WAP enabled phones. Although the screen real-estate of these devices is still very small, most of them come equipped with high-resolution color interfaces and predictive text input capabilities - a significant improvement over 1st generation WAP-handsets. The second class of device includes next-generation smart phones with 3G capabilities. The majority of these devices encompass larger screens, higher-resolution interfaces, enhanced Web-browser support and multiple input options. For example, most of these phones support XHTML, (e.g. Nokia 6630, Siemens SX1) and some of the phones incorporate flip-out messaging keyboards, (e.g. Nokia 6820, Siemens SK65). The final class of mobile device includes PDA's, XDA's and other small handhelds. These devices have much larger screens, higher-resolution interfaces, integrated Web-browsers supporting standard HTML as well as a multitude of input options including full messaging keyboards and

stylus. So while previous incarnations of the Mobile Internet were not viewed as a success there is increasing optimism about the future of this important information medium. These new handsets, coupled with improved content, more attractive billing arrangements and greater bandwidth appear to be attracting more and more users to the Mobile Internet on a daily basis. This in turn is likely to lead to a need for improved search facilities, facilities that we argue need to be carefully adapted to the features of the Mobile Internet, its devices and the differing needs of its users

2 Characteristics of Mobile Search

We do acknowledge that even current mobile devices suffer from inherent limitations that can greatly effect the usability of many mobile applications [Ramsay and Nielsen, 2000]. This is largely due to the fact that very little effort has been made to cater for the unique requirements of the Mobile Internet. Such limitations can be especially problematic for mobile search applications. Small displays restrict the number of results that can be displayed on a single screen and input limitations impact the type of queries that are provided. Most mobile search engines offer only simple adaptations to standard Web search interfaces. For example, fewer results are returned per page and the *snippet* text associated with each result may be truncated or simply removed [Jones *et al.*, 2002]. We believe that more fundamental adaptations are necessary.

In particular we believe that displaying snippet text alongside each search result is not feasible in the mobile realm. However we also think that displaying just a title or a url to represent each search result does not provide the user with meaningful information about the context of a given result. Instead we propose the use of *related queries* in place of snippet text as a more economical form of result gisting. Economy of screen-space is a major concern in mobile search and the use of related queries can significantly increase the number of results displayed per screen on mobile devices. Moreover, in this paper we also show that the terms in these related queries can be more informative, when it comes to capturing the content of an associated page, than the page's snippet text. Similar work has also been carried out in [Jones *et al.*, 2004], where key phrases, automatically extracted from web pages are used to provide users with accurate, meaningful representations of search results on mobile devices.

3 Towards Intelligent Result Gisting

A personalized search engine called I-SPY, [Smyth *et al.*, 2003], can provide these related queries by tracking and recording past queries that have resulted in the selection of a given result page. I-SPY is a community-based meta search engine that records the search histories of communities of like-minded users [Smyth *et al.*, 2003; Forthcoming]. Each time a user selects a result, p , for a query, q , I-SPY updates a counter in a *hit-matrix* to reflect this selection. This information is used to adapt future result-lists for similar queries by promoting results that have been selected in the past.

One of I-SPY's distinguishing features is that it can recommend related queries alongside a search result. For each search result p , that is associated with some target query q , I-SPY can generate a set of related queries from those queries, q_1, \dots, q_n , that have also led to the selection of p [Smyth *et al.*, Forthcoming]. These related queries can also be ranked according to the likelihood that they match a user's current requirements. These ranking strategies are described in [Balfe and Smyth,]. Figure 1 illustrates these related queries in action. If we take, for example, the query 'MP3' we can see that the first search result, 'mp3.com' is associated with queries such as 'music' and 'charts'. These *related queries* can inform the user as to the context of the search result in question.




Figure 1: Query 'MP3' on a mobile phone

It is interesting to consider whether these related queries could be used as an alternative to snippet text to gist search results, especially given the interface limitations of mobile devices. Preliminary empirical evidence generated from a similar evaluation we carried out suggests that these related queries may be as informative as snippet text [Church *et al.*, 2005]. In our previous evaluation we used a smaller dataset and a different independent search engine, namely Google. In order to validate our earlier findings, we decided to carry out a more in-depth evaluation using a different independent search engine, namely HotBot and a larger dataset. The details of this evaluation are described in the following section.

4 Evaluation

The data for this evaluation was collected by mining the search logs of I-SPY. Specifically we collected a set of 684 result-pages that were selected by searchers in response to various queries. Each of these pages was associated with at least 2 related queries and a unit of snippet text generated by Google. The purpose of our experiment was to consider how well these related queries represented the page in question, relative to its snippet text. To do this we supposed that

the representativeness of a set of terms relative to some page could be measured by the position of the page in the result-list generated by some search engine when using these terms as a query. Hence, in our evaluation we transformed each set of related queries and piece of snippet text into a new search engine query, submitted the queries to HotBot and compared the position of the target page, p , in the two result-lists produced. We compared a number of variations of this approach by constructing our search queries in different ways.

4.1 Query Generation

Probably the most important part of our evaluation was the generation of the test queries from both the related queries and snippet text. In all, six query generation strategies were tested: two that produced queries from the terms contained in the related queries for a page and four that used terms from the page's snippet text. Strategy *RQ1* produces a test query by concatenating the related query terms into a single query. Strategy *RQ2* uses a similar approach but duplicate terms are removed.

The snippet text conversion was slightly more complicated. To ensure a fair evaluation, we generated a test query from the snippet text using the same number of terms as the test queries produced from the related queries. To generate the test queries we parsed the snippet text to remove stop-words and special characters and then selected terms from the remaining snippet text using four different strategies. In strategy *S1* we select a random set of k terms, where k is the number of terms in the test query produced by *RQ1*. Strategy *S2* selects the top k most common terms in the snippet text, where k is the number of terms in the test query produced by *RQ2*. For strategy *S3* we again select a random set of k terms but this time k is the number of terms produced by *RQ2*. Finally, strategy *S4* selects the top k most common terms in the snippet text, where k is the number of terms produced by *RQ1*.

4.2 Relevance Assessment

After submitting each test query to the HotBot search engine, we examined the top 500 HotBot results only and compared the position of p , the target result-page, in the result-lists produced for each test query. The higher p is in the result-list the more representative the test query must be as an indicator of p 's content and hence the more representative the related queries or snippet text. Along with this positional information we also examined the percentage of results matched by each test query strategy as well as the average length in number of terms of the related queries vs. snippet-based strategies.

4.3 Results

First we looked at the average position of each p in the result-lists produced by HotBot, see Figure 2. When calculating the average position, a penalty is applied when the target result cannot be found for the query in question. If the target result cannot be found, p is given a default position value of 501 (because we are only examining the first 500 HotBot results). Both related query strategies, *RQ1* and *RQ2*, perform very well. *RQ2* locates p at an average position of 39 in the result-

list compared to the best performing snippet-based strategy which locates p at an average position of 103.




Figure 2: Average position of p in result-lists

The poor showing for the snippet text strategies could have been due to frequent penalties being incurred when p was not present in the top 500 results. To understand this we examined the percentage of results found for each test query strategy in the top 500 HotBot results, see Figure 3. The related query strategies succeed in producing result-lists that contain p for between 92% and 94% of queries, a significant improvement when compared to the snippet text strategies which return p for between 66% and 80% of the test queries.




Figure 3: Percentage found in the top 500 HotBot results

These results suggest that the terms contained in the related queries are more representative of the pages they refer to than an equivalent number of terms taken from the snippet text associated with these pages. And by using related queries instead of snippet text we can achieve a significant saving in display-space. For example, as shown in Figure 4, our related queries contained only 4 unique terms on average, compared to snippet texts with an average of 35 terms (words), or 21 terms with stop-words removed.

5 Conclusion

One of the main issues with mobile search engines concerns the manner in which search results are displayed. In this paper we have suggested using related queries as an economical alternative to snippet text for result gisting. Our evaluation suggests that the terms in related queries have the potential to better capture the essence of their associated pages than the terms in the snippet texts, and so may serve to be a more




Figure 4: Average number of terms in test queries

informative gisting approach. Moreover, related query terms take up a small fraction of the screen space associated with the display of snippet text, which is a major advantage for the provision of mobile search.

6 Acknowledgements

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

References

- [Balfe and Smyth,] E. Balfe and B. Smyth. Collaborative Query Recommendation for Web Search. In *Proceedings of 16th European Conference on Artificial Intelligence*, pages 268–272, Valencia, Spain.
- [CellularOnline, 2004] CellularOnline. See, <http://www.cellular.co.za/>, 2004.
- [Church *et al.*, 2005] K. Church, M. T. Keane, and B. Smyth. An Evaluation of Gisting in Mobile Search. In *Proceedings of the 27th European Conference on Information Retrieval*, Santiago de Compostela, Spain, 2005.
- [Jones *et al.*, 2002] M. Jones, G. Buchanan, and H. Thimbleby. Sorting out Searching on Small Screen Device. In *Proceedings of the 4th International Symposium on Human Computer Interaction with Mobile Devices*, Pisa, Italy, 2002.
- [Jones *et al.*, 2004] S. Jones, M. Jones, and S. Deo. Using Keyphrases as Search Result Surrogates on Small Screen Devices. *International Journal of Personal and Ubiquitous Computing*, 8(1):55–68, 2004.
- [Ramsay and Nielsen, 2000] M. Ramsay and J. Nielsen. The WAP Usability Report, 2000.
- [Smyth *et al.*, 2003] B. Smyth, E. Balfe, P. Briggs, M. Coyle, and J. Freyne. Collaborative Web Search. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 1417–1419, Acapulco, Mexico, 2003.
- [Smyth *et al.*, Forthcoming] B. Smyth, E. Balfe, J. Freyne, P. Briggs, M. Coyle, and O. Boydell. Exploiting Query Repetition & Regularity in an Adaptive Community-Based Web Search Engine. *User Modeling and User Adapted Interaction*, Forthcoming.