The Role of Context in Information Retrieval

Gareth J. F. Jones
School of Computing &
Centre for Digital Video Processing
Dublin City University
Glasnevin, Dublin 9, Ireland
Gareth.Jones@computing.dcu.ie

Peter J. Brown
Department of Computer Science
University of Exeter
Exeter EX4 4QF, United Kingdom
P.J.Brown@exeter.ac.uk

ABSTRACT

The scope for information retrieval applications continues to growth to encompass larger and more diverse archives and new computing environments. The consequent increased demands on information retrieval systems continue to motivate research into new and more efficient search technologies. Users of information retrieval systems work in personal and physical contexts, while the documents containing the information they seek often relate to specific contexts. It is argued that taking account of context can improve information retrieval effectiveness. While potential use of context has certainly been under explored, it is already present in established techniques such as relevance feedback. The emergence of new information retrieval environments, such as those associated with mobile computing, raises new challenges for information retrieval for which greater use of context may form a important component. Among the many questions raised by attempting to perform information retrieval in context are issues of how algorithms might be extended or developed to facilitate use of context, does the user need to be actively involved to make use of context, where should the context information come from, and how might the effectiveness of such methods be tested in extended evaluation frameworks?

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval

General Terms

Algorithms, Experimentation

Keywords

context, information retrieval

.

1. INTRODUCTION

Information retrieval (IR) is often defined in terms of the location and delivery of documents to a user to satisfy their information need. This apparently simple task turns out to be highly complex. In order to facilitate research and move towards real systems various pragmatic assumptions and simplifications have traditionally been made in IR research. A key factor that has generally been factored out is explicit modelling of the *context* in which the search is being carried out, although as we demonstrate later it is already used implicity in at least one IR technique.

The notion of context has been one of the mostly widely interpreted, or possibly abused [1], terms in IR, encompassing topics such as interaction, domain specific search, personalization, [2] and environmental sensors [3]. None of these or other interpretations are necessarily wrong, but neither we feel are they exclusive.

IR has conventionally been concerned with users working inside an "information world" using desk-based systems. However, developments in mobile and wireless computing mean that users of IR systems may now be embedded in the "real/physical world" surrounded by environmental sensors providing rich information of physical context [4].

An important question then is, how should "information retrieval in context (IRiX)" be interpreted? Should it refer only to taking account of the information world, and within this refer only to what has often classified as "interactive IR"; should the scope be broader encompassing personalization and implicit feedback [2], or should it be broader still to covering mobile IR as well [4]. Further than this, from our analysis, once one begins to look for issues of context within IR it starts to emerge everywhere. Should IRiX then focus on all forms of context information wherever it can be found in the IR process? One interesting question with regard to IR and context that we ask ourselves is, does IRiX really imply interaction in IR? Are there opportunities for context to be exploited without any user involvement or modelling of the individual user, we argue that the answer appears to be yes.

However the scope of IRiX is defined, the overall aim in taking account of context within IR is to improve the efficiency with which the user's information need is satisfied. The appropriate measures of effectiveness here will themselves often depend on the informational or physical context in which the search is being conducted. This may perhaps imply that the user has to do less work to satisfy their information need if context is incorporated into the search process, although depending on how the context data is en-

tered and used, it may actually be harder work for the user. By some measures the complexity of the retrieval process might be reduced by incorporating context, but by others it may well be increased.

The remainder of this paper analyses the use of context within current IR, discusses the challenges and opportunities of IR for mobile and ubiquitous computing environments, and outlines some possibilities for exploring the extension of the role of context within IR.

2. CONTEXT IN CURRENT IR SYSTEMS

While it is to a large extent true that current IR systems take no account of context explicitly, any information given to the IR system by the user and used in computation of the output can be regarded as some form of search context.

Current IR models generally assume that document attributes, typically words, are independent [5]. This is a practical approximation enabling the development of computationally tractable algorithms with parameters that can be reasonably estimated from the data available. However, even within this model we can think of attributes appearing in the context of other ones. For example, relevance feedback (RF), already implicity makes use of attribute context. Within RF documents are placed in the context of their relevance to information need rather than seen as independent items within an overall document collection [6]. Feedback from the user provides context dependent information to the IR system on the relevance of a specific document. Using this feedback information, RF seeks to model the importance of terms to the information need and expands the request to include those terms deemed to make it a better expression of the information need. Interestingly pseudo or blind RF, which operates without relevance input from the user effectively, automatically develops context between documents retrieved using the initial search. Search context is being inferred here without any user involvement. Some RF methods take account of word context within documents, for example local context analysis [7] and document summaries [8].

IRiX for searching in the information world relates to taking account of the user and their cognitive attributes. This may result in the search being personalized or the user placed within a user class [2]. While it is sometimes difficult even to elicit user feedback of document relevance, personalization of search often seems to be based on elaborate interfaces requiring considerable engagement from the user [9] or marking relevant material for feedback [1]. While these interfaces may be effective when used properly, even by novice users, it must surely be open to question whether the typical user with a real information need will make the effort to use them properly. Perhaps we should focus our attention for IRiX more within the system. As the next section illustrates this is an even more significant question for mobile IR.

3. CONTEXT IN MOBILE AND UBIQUITOUS INFORMATION ACCESS

Mobile computing devices range from "semi" portable, e.g. a laptop, to truly mobile, e.g. mobile phones and PDAs, indeed sometimes the device may effectively be embedded in the user's environment, e.g. a car computer. This change in the working environment has many implications for the user needs in respect of IR which involve information and context.

Various studies have considered the implications of this new environment for IR [10][4][11][3]. Some important conclusions of these and other studies are that:

- the user is often engaged primarily in another activity other than information searching and is likely to want rapid access to relevant information to assist them.
- the user is embedded in the "real world", and will often require information accurate up-to-date information relating to their current physical circumstances.
- particularly for small handheld devices anything more than cursory searching and browsing is impractical.
- the potential of physical context data to improve IR precision appears attractive and should be explored.

Context in these studies has tended to focus on the physical world, rather than the information world more typically associated with interactive IR. The issues of information world IRiX still apply here, although they may be modified with respect to the types of user activities or searching. Here we consider only the additional context features associated with the physical world.

In mobile IR applications much physical context data is potentially available via personal and environmental sensors; for example, the user's location, who they are with, ambient temperature, traffic conditions, nearby shops and offices, and their attributes. Some of this information might be used directly, while other data must aggregated to infer higher level context or user activity before it is useful in IR. As well as its use in IR, this context data might also be used to determine the best mode of information delivery, e.g. via audio output while the user is driving [3].

In previous work we have explored issues of IRiX relating to the physical world for mobile computing environments. In this work we developed the concept of the *context-of-interest* [11]. The basis of this proposal is that mobile users are often likely to be most interested in information referring to their future context rather than their current one, which may be out-of-date when the information is actually delivered. This led us to consider the topic of *prediction* of physical context. This may be based on user independent physical context tracking or be personalised through incorporation of information from a user's *context diary* of recent, current and expected activities.

It has also been noted that users of small mobile devices are less able to interact with them [10][12]. Thus, IR systems on these devices must behave more autonomously than desk-based systems, and attempt to provide users with the relevant information they need without significant browsing.

Also as noted above users may often be engaged in other activities rendering them unable to search for information of use to them in their current physical context, or they may not be aware of the existence of information which they might find useful. In this situation it would be useful if the IR system were to behave *pro-actively* to look for potentially interesting information by making use of context data to form form search queries automatically. This is more akin to information filtering or routing than IR, but is nevertheless a further example of non-interactive IRiX.

4. CONTEXT IN IR ALGORITHMS

While the physical operating environment and means of gathering context information are different for traditional desk-based IR and mobile IR, the issues of how to use context within IR algorithms raise fundamentally the same questions, and may have the same or related solutions.

A key concern is how to model context data as attributes within an IR system, and how to incorporate these within IR algorithms. This raises questions of establishing context classes and defining context attributes within these classes. Linguistic context features can be sometimes be handled in a straightforward way using relevance feedback type algorithms for learning and personalization, although handling even these features will often be more complex than this. More problematic in general is other context data for which the attributes are not obvious and for which suitable similarity coefficients must be sought.

One simple way to use context in IR would be to use them to introduce boolean type constraints on existing search mechanisms, e.g. to limit the time range searched for documents. This is a very limited vision of the use of context. In some situations there may be a balance between the level of context match and the level of content match. A document which is highly matched based on context while poorly matched based on content, may be judged by the user to be more or less useful than one which matches poorly on context and highly on content [13]. Matching algorithms for IR in context should be able to take account of this.

Consider the popular context features of time and location. These have continuous values that can be described in precise or vague terms at different levels of granularity. For example, an event may take place at `9:57`, at `about 10 o'clock`, `in the morning`, and at a given numerical grid reference, or in a named street, district, or city. In this case it is not clear what the appropriate representation is and how the significance of each feature should be defined in terms of selectivity. Determining how to compute a scored match between context fields of varied specificities, and how to integrate these scores with traditional IR content-based query-document matching in a well-founded way is a key issue here. A novel approach to matching between temporal and location fields for topic detection and tracking is introduced in [14]. These methods assign a score to degree of temporal overlap between an event in two documents and compute a matching score between locations using a geographical ontology tree, e.g. `Paris` is in `France` which is in `Europe`. While interesting, these methods do not currently improve retrieval effectiveness beyond standard content matching. A further important issue is the extent to which the context associated with documents and queries has to be stated explicitly or can be extracted automatically from their contents [2][3]. Finding potential context information within a document is only part of the problem, it may then be necessary to try to find associations between context data using shallow natural language processing methods, perhaps in a manner similar to event construction in information extraction methods as applied in question-answering systems.

Once some similarity measure between context of the same type has been established this could be incorporated in an extended form relevance feedback. Related context information in retrieved documents could be identified and used to extend the contextual annotation of the individual documents, to determine the likely importance of this context information across the document collection, and then to extend the search request to include significant context labels. The notion of using cross collection statistics to establish reliable features associated with an event has been already been in explored in [15], and this might be extended to determine significant and reliable context labels. The success of this idea requires the development and extension of a number of techniques, but could offer a way to integrate context of various forms into an IR search.

5. REFERENCES

[1] L. Finkelstein, et al. Placing Search in Context: The Concept Revisited. *ACM Transactions on Information Systems*, 20(1):116–131, 2002.

[2] S. Lawrence. Context in Web Search. *IEEE Data Engineering Bulletin*, 23(3):25–32, 2000.

[3] G. J. F. Jones and P. J. Brown. *Context-Aware Retrieval for Ubiquitous Computing Environments*. In *Proceedings of the Mobile HCI 2003 International Workshop on Mobile and Ubiquitous Information Access*, pages 227–243, Udine, 2003.

[4] P. Coppola, V. Della Mea, L. Di Gaspero and S. Mizzaro. *The Concept of Relevance in Mobile and Ubiquitous Information Access*. In *Proceedings of the Mobile HCI 2003 International Workshop on Mobile and Ubiquitous Information Access*, pages 1–10, Udine, 2003.

[5] C. J. van Rijsbergen. *Information Retrieval* (2nd edition), Butterworths, 1979.

[6] G. Salton and C. Buckley. Improving Retrieval Performance by Relevance Feedback. *J. American Society of Information Science*, 41(4):288–297, 1990.

[7] J. Xu and W. B. Croft. Improving the Effectiveness of Information Retrieval with Local Context Analysis. *ACM Transactions on Information Systems*, 18(1):79–112, 2000.

[8] A. M. Lam-Adesina and G. J. F. Jones. Applying Summarization Techniques for Term Selection in Relevance Feedback. In *Proceedings of ACM SIGIR 2001*, pages 1–9, New Orleans, 2001. ACM.

[9] J. Pitkow, et al. Personalized Search. *Communications of the ACM*, 45(9):50–55, 2002.

[10] B. J. Rhodes and P. Maes. Just-In-Time Information Retrieval Agents. *IBM Systems Journal*, 39(3&4):685–704, 2000.

[11] P. J. Brown and G. J. F. Jones. *Exploiting Contextual Change in Context-Aware Retrieval*. In *Proceedings of the 17th ACM Symposium on Applied Computing (SAC 2002)*, pages 650–656, Madrid, 2002.

[12] C. Gurrin, et al. *Mobile Access to the Fischlár-News Archive*. In *Proceedings of the Mobile HCI 2003 International Workshop on Mobile and Ubiquitous Information Access*, pages 124–142, Udine, 2003.

[13] G. J. F. Jones and P. J. Brown. *Information Access for Context-Aware Applications*. In *Proceedings of the ACM SIGIR 2000*, pages 382–384, Athens, 2000.

[14] J. Makkonen, H. Ahonen-Myka and M. Salmenkivi. Simple Semantics in Topic Detection and Tracking. *Information Retrieval*, 7(3&4):347–368, 2004.

[15] S. Dumais et al. Web Question Answering: Is More Always Better? In *Proceedings ACM SIGIR 2002*, pages 291–298, Tampere, 2002. ACM.