Journal of Information Retrieval, 2, 1–5 (1992)

Incorporating aspects of information use into relevance feedback

IAN RUTHVEN
Department of Computing Science, University of Glasgow, Glasgow, G12 8QQ, Scotland

Received August 3, 1998; Revised April 8, 1999

Editors: M. Dunlop and M. Lalmas

Abstract. In this paper we look at some of the problems in interacting with best-match retrieval systems. In particular, we examine the areas of interaction, some investigations of the complexity and breadth of interaction and attempts to categorise user's information seeking behaviour. We suggest that one of the difficulties of traditional IR systems in supporting information seeking is the way the information content of documents is represented. We discuss an alternative representation, based on *how* information is *used* within documents.

Keywords: relevance feedback, interaction, information seeking

1. Introduction

Modern information retrieval (IR) is an inherently interactive activity and users can expect to engage in a variety of tasks and techniques in the course of an information-seeking session. The IR system should support a number of operations, e.g. learning about the system, learning about the information source, describing an information need and giving feedback.

Each of these areas require users to make decisions about what actions to take at a particular stage. What actions the user performs depend not only on what the system does but what the user thinks it does. For example, Cool et al [4] indicate that users of boolean systems can successfully use existing strategies on ranked retrieval systems but this success depends highly on their mental model of how the system operates. Other authors, e.g. [7], have investigated more general information seeking techniques that are not directly dependent on the IR system being used.

How to provide effective support for the interaction between a user and an IR system raises a number of difficult questions, partly in defining appropriate methods for individual tasks, but also more general questions about the interaction process itself.

For example, we may need to consider the balance between the amount of interaction involved in particular actions, e.g. should we try to minimise the time spent formulating and reformulating queries but with improved methods of making relevance assessments? Should a user with an ill-defined information need and who is unfamiliar with a particular system expect to put in more effort to find information or should the system take a more controlled approach by directing the user towards information?

As well as balance of interaction we should also consider the complexity of the interaction. Given that users will need support throughout a variety of types of search (well-defined, ill-defined, searching for a known object, etc.) do we want multiple methods of requesting information for different types of search or does this impose too large a learning burden on the users?

A third area is the integration of the IR system and the environment in which the search is being performed. To what degree do IR systems need to incorporate external factors (context, task, user knowledge) into a search? Do these situational factors improve the individual experience of searching and can we include enough of this information to make a difference?

All these issues (and many more) have an impact on how successfully individual users and types of users (novice, experienced with IR systems, experienced with a particular system) can use IR systems. In addition they raise the question of whether we can expect a single IR system to support all users in all searches. This is the concern of the next section which examines some of the calls for improved interaction techniques in information retrieval. This is followed by a description of two particular difficulties that arise with current IR systems. The paper concludes with a proposal for an alternative method of representing how information appears in documents in order to improve the quality of the users interaction.

2. Modelling the interaction

In [3], Belkin et al. consider the types of search a system must support. They start out from the observation that people engage in multiple information seeking behaviours, both within and across a search. Thus if different types of information seeking (searching by query, browsing, etc.) are to be supported effectively, IR system design and implementation must be capable of adapting to the users changing requirements of the system. They present a classification of 16 types of behaviours. Each type is defined by its values on each of four dimensions (*method of searching* - either scanning or searching, *mode of retrieval* - recognising a relevant item or specifying it, *goal of retrieval* - learning about the system/collection or selecting relevant information and the *resource considered* - information items or meta-information). Each type of information seeking behaviour may also be regarded as an IR system in itself.

What motivates the choice between one type of search and another may be personal searching style but it could also arise from the *uncertainty* at various points in the search. For example, a user may choose to recognise relevant items rather than specify his information need if he is unsure of what information is available, or he choose to use meta-information, such as a thesaurus, if he is unsure of the vocabulary used in the collection. Kuhlthau, [10], promotes uncertainty of searching as the primary factor in characterising the move from one stage in a search to another and the particular choice of activity.

Bates, [1], on the other hand considers the complexity of interaction. She presents a classification of activities based on the conceptual complexity of user actions (ranging from move, individual actions such as marking a document relevant, to strategy, a plan for a whole information search). This categorisation allows her to speculate on how IR could support interaction with the rather bleak observation that most, current IR systems only support very limited interaction with very little support for more complex information seeking patterns.

The main reason for this is that IR systems generally only support discrete actions, e.g. selecting query terms, marking documents as relevant, etc., rather than more complex information seeking behaviours. However the problem in supporting more difficult interaction is not simply due to a lack of operators or features to perform them but also a lack of information about the system and control on the users part.

Authors such as Ingwersen, [8], have stressed the importance of taking into account situational factors such as the task the user is trying to perform, what the user knows about the domain, system, their cognitive environment and the conceptual aspects of searching. One possible advance presented is the use of an automated intermediary which encapsulates knowledge about the system, domain, user, etc., based on empirical findings of searching and expert analysis of searching behaviour and mediation. This intermediary with appropriate dialogue techniques acts as an expert suggesting techniques or options to a user and mediating between the user and system. We must however consider how much knowledge of searching is necessary to build such models, how to represent this information and how to update it. Nevertheless the possibility that the system may suggest techniques or provide an expert aid may prove to be a powerful tool in IR.

Each of these studies, and other behavioural studies such as [6], demonstrate how complex the users actions may be, both in what they want to do and what factors affect their choice of action. The next section discusses two aspects of current IR systems that can influence the complexity of these actions.

3. Feedback and control in interaction

Two characteristics of IR systems that can make them difficult to use are the lack of *feedback* or explanation provided by the system as to what it is doing, and the lack of *control* users have over their intended actions.

IR systems expect users to make decisions throughout an IR session but typically do not provide enough information for them to decide what factors the decision should be based on. For example in interactive query expansion, e.g. [2], systems ask users to select terms useful in retrieving relevant documents, but do not necessarily give the users either any indication of the context in which these terms may appear or of how important this term will be regarded in relation to other query terms. A similar scenario occurs in relevance feedback. How many documents should a user select to improve the retrieval? What documents should they be selecting - all those that are relevant in some way, all those that contain extremely relevant information, all those that contain a lot of relevant information and little irrelevant information? This type of choice can have a significant affect on the retrieval performance but how to make the choice is not clear.

Uers cannot predict the effect of their interaction. For example when users select query terms they have little indication of how each term in the query will affect the result, even worse when they start to combine terms. This lack of indication of how users should make choices, such as what query terms to put and what documents to mark as relevant, mean that users can be hampered by being forced to make a decision without knowledge of the effect that decision is going to have on their search, and what outcomes could have been achieved by alternative actions. Increased interaction means increased control and increased responsibility but not necessarily increased information on how to use this control.

Some authors such as [13], propose increasing the user's awareness of the variability of the IR system functionality by increasing the clarity of instructions at the interface but, although this may reduce the uncertainty about how to express a query, it does not reduce the uncertainty of searching as a whole. The aim of this work-in-progress is to provide more precise ways of interacting with an IR system by tackling the representation of information and interaction.

4. Representing the interaction

Part of the reason why the difficulties of feedback and control arise in IR is that the design of systems concentrate on the matching function; interaction is geared towards providing better inputs to this function rather than ways of manipulating information objects. Even techniques such as interactive query expansion which aim to give the users more control and produce more interaction do not necessarily result in the users interacting more meaningfully with the system. In [5], Denos et al make the good point that although users can make explicit relevance judgements on why a document may be relevant to an information need, current systems have little means of using this information.

In IR, representations of information are defined for a single purpose - retrieval - but the representation of the interaction needs to support the way information is used and managed by users. In addition, this representation must allow the users a better understanding of what the systems do and how decisions are made. Reducing the conceptual gulf between what systems do and what users think they may be doing also allow users better control over their search.

Traditional IR systems use a procedural approach to representation, specifying how keywords, phrases, etc should be selected for retrieval and how they should be prioritised. We are arguing for a declarative approach to representation. Rather than representing the information available or how important it may be, we need to represent the way information is used. Declarative representations can be used to express propositions about terms, making explicit the characteristics of a term as it appears in a document, e.g. the local context in which a term appears, or its topical nature within a document. A term is not described simply by its absence/presence and frequency within a document but also by statements about its role in the document.

This does not mean throwing away successful techniques for information representation but using them in a more approachable way. For example, although successful techniques have been developed using e.g. contextual information or topical structure or discourse analysis, they tend to be used automatically across a collection rather than when required or in response to evidence from the user. The aim is not to make specification of an information need more complicated *a priori*, although it would be possible to query by characteristic e.g. specifying the main topic of the document, but to use characteristics of information use to explain decisions to the user and to give users more power to develop their own strategies, i.e. displaying why documents are retrieved and allowing users to include/exclude criteria for retrieval.

By expanding the representation of a term, relevance feedback, [9], should not simply be a process of selecting good terms for retrieval but a two-step process: detecting characteristics of relevance, and adapting the search process itself. In the first step we can analyse documents to detect why a document may have been relevant, making more use of the relevance information given by user. For example,

- there may be a large overlap in the content of marked relevant documents but not the structure, or local context in which terms appear. This may correspond to a search in which the user wants all possible information on a topic.
- there may be a high structural similarity in the relevant documents, e.g. the same terms are related to the main topic of the document but appear in a variety of contexts. This may be a search for information that the collection has a lot of information on (so the user only wants whole documents about the subject) or the user is finding his way around the collection.
- or there may be a high similarity in context, but not in content or topical relationships. This may corresponds to searches where information is only relevant in certain combinations, for example `'information retrieval systems'` not `'information retrieval'` or `'information...systems'`. '
- relations between terms, such as causal relations e.g. 'smoking causes cancer, can also indicate why a particular document is relevant

The second step suggests that retrieval is not simply a matching process or a decision process about what to retrieve but on how to retrieve. This step would be to derive a strategy for retrieving more documents such as a selective combination of evidence.

In [12] we demonstrated that incorporating information on how are used within documents - term characteristics, in a RF situation, can lead to significant improvements in retrieval effectiveness across collections. We also demonstrated, experimentally, that different combinations of characteristics are more suitable for different queries. In other words, different combinations of characteristics are better at detecting relevance for different queries.

This information was used in [12] to motivate biasing retrieval in favour of how a term was used in relevant, as opposed to irrelevant documents. In our experiments on the TREC collections we demonstrated that this technique of selective relevance feedback - selecting which characteristics to use for which query - not only performed well but outperformed standard relevance feedback algorithms such as the $F_{4.5}$ measure, [11].

This can be extended by analysing, on an iteration-to-iteration basis, how to combine evidence. By taking account of the differences between iterations, the temporal nature of a search, we can also connect the search to how to display the information, e.g. displaying how the systems view of search is changing. We can also suggest possible alternative actions to the user such as displaying relationships between documents.

Part of the difficulty Bates [1], found in supporting searching was that users want to know what information is being introduced/rejected and why. We cannot predict what information users will want to use or what information they want, so a better explanation of how current systems operate, e.g. do we use stop words, how is a term stemmed, etc. are not necessarily going to improve user's satisfaction unless we examine how users search and what criteria they need. Simply explaining how we do retrieval is not necessarily a better option, unless, as here, we related them directly to what the users are interacting with. More complex search manipulation techniques such as those described in [1] need better support by better descriptions of what the users are manipulating, rather than by simply more sophisticated matching algorithms.

5. Conclusions

In this short paper we argue that current IR systems do not adequately support users; in particular by not providing meaningful feedback on the system's functioning or enough context to control a search. This arises in part from the gulf between how the system asks for information, how it translates the user's intent into system questions, and how it uses this information. A more approachable option is to present more meaningful descriptions of information objects based on the characteristics of terms within documents and how they are used to retrieve objects. This allows users to access systems function, develop their own strategies and support better information manipulation within an information retrieval session.

6. Acknowledgements

We would like to thank Mounia Lalmas for her valuable comments on this paper and also the anonymous referees for their comments and suggestions.

References

1. M. Bates. Where should the person stop and the information search interface start? *Information, Processing and Management*, 26(5):575–591, 1990.
2. M. Beaulieu. Experiments with interfaces to support query expansion. *Journal of Documentation*, 1(53), 1997.
3. N. J. Belkin, C. Cool, A. Stein, and U. Thiel. Cases, scripts and information-seeking strategies: On the design of interactive information retrieval systems. *Expert Systems with Applications*, 9(3):379–395, 1995.
4. N. J. Belkin J. Koenemann C. Cool, S. Park and K. B. Ng. Information seeking behavior in new searching environment. In *CoLIS 2, Proceedings of the Second International Conference on Conceptions of Library and Information Science: Integration in Perspective*, pages 403–416, Royal School of Librarianship, Copenhagen, 1996.
5. N. Denos, C. Berrut, and M. Mechkour. An image system based on the visualization of system relevance via documents. In *Database and Expert Systems Applications (DEXA 97), 8th International Conference*, Toulouse, France, 1997.
6. D. Ellis. A behavioural approach to information system design. *Journal of Documentation*, 45(3):171–212, 1989.
7. V. Florance and G. Marchionini. Information processing in the context of medical care. In *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 158 –163, Seattle, Washington, USA, 1995.
8. P. Ingwersen. *Information Retrieval Interaction.* Taylor-Graham, 1992.
9. J. J. Rocchio Jr. *Relevance feedback in information retrieval*, chapter 14. Prentice-Hall, 1971.
10. C.C. Kuhlthau. Principle for uncertainty for information seeking. *Journal of Documentation*, (4):339–355, 1993.
11. S. E. Robertson and K. Sparck Jones. Relevance weighting of search terms. *Journal of the American Society of Information Science*, 27:129–146, 1976.
12. I. Ruthven and M. Lalmas. Selective relevance feedback using term characteristics. In *CoLIS 3*, 1999. to appear.
13. B. Shneiderman, D. Byrd, and W. B. Croft. Sorting out searching: A user-interface framework for text searches. *CACM*, 41(4):95–98, April.