Abduction, explanation and relevance feedback

Ian Ruthven
Department of Computing Science
University of Glasgow
G12 8QQ
Email: *igr@dcs.gla.ac.uk*

1. Relevance feedback

Relevance feedback (RF) is an interactive technique that is intended to automatically improve an information retrieval (IR) system's representation of a search based on documents that a user has assessed as being relevant. One strength of RF is that it requires minimal input from a user: a user only has to *indicate* relevant material - they do not have to describe what makes a document relevant.

The majority of existing RF techniques produce a new query based on what terms appear in relevant and irrelevant documents; they do not consider how terms are *used* in documents. A document may only be relevant if the terms appear in a certain context, if specific combinations of terms appear, or if the main topic of the document is important. Therefore we need to be able to detect *why* a document might be relevant.

Our research centres around defining a formal model of RF that is based on detecting what aspects of a term's use in a document - *term characteristics* - are good indicators of relevant material. This model also incorporates behavioural aspects of searching in order to offer better support to *how* users search.

2. Term characteristics

Term characteristics give values to terms according to how they are used within documents or collections. Examples include standard weighting schemes such as *idf*, *tf*, and can represent aspects of term use such as local context, causal relations, and relation to the main topic of the document.

In previous work, [RL99], we demonstrated that using relevance assessments to *select* which characteristics of a term's use should be used to modify a user's query can significantly improve retrieval effectiveness. However, selecting which characteristics are good indicators of relevance for each term is subject to a number of variables

For example, different documents may be relevant to different parts of a user's information need so one combination of characteristics may be good at retrieving some relevant documents and poor at retrieving others. In addition, documents may be relevant for many reasons and different characteristics may be better or worse at retrieving different relevant documents for a query.

In [RB99] we showed that how the user searched for information was also important in deciding how the selection procedure should operate. For example, different methods of selecting term characteristics worked better if the user was only looking for highly relevant documents than if they were searching for documents that contained any relevant information.

In order to deal with the diversity of system and user variables that can affect how term characteristics should be selected we are formalising a model of RF that is based on *abductive inference* or *abduction*.

3. Abduction and explanation

Abduction is primarily a model of explanation. Given a set of relevance documents, an abductive inference system will generate a set of possible explanations for why the documents may have been marked relevant. Each of these explanations will be a set of weighted characteristics of terms. Not all explanations are equally good at explaining the relevant assessments; in order to choose the optimal or best explanation we need to infer information from three sources: the user behaviour, the content of the relevant documents, and the previous search history.

3.1 Inference from user behaviour - inferring state of search

How the user assesses documents can be important in detecting both what the user finds relevant and identifying stages in a search. Evidence from the use of partial relevance assessments can help in identifying when a search is focused and evidence from the consistency of searches (based on the document content) can help in identifying searches for which the relevant material is distributed between documents.

Thus we can choose, at each iteration, which of the relevant documents we want to consider: - all the relevant ones, only the highly relevant ones, the most consistent ones, or the ones that we feel may have been more central to the user's relevance assessments.

3.2 Inference from content of relevant documents

The selection of good explanations in part depends on the nature of the characteristics that compose the explanations. Potentially any set of weighted characteristics of terms can provide an explanation, however we

can cut down this search space in a number of ways, incorporating standard RF and IR theory. For example we can assume that a term that is not in a relevant document is not a good indicator of relevance at the current search stage. A good term or characteristic should also discriminate well between relevant and irrelevant documents so we should not consider any explanation in which a term or characteristic is better expressed in irrelevant documents than relevant ones. This inference stage will select good components of possible explanations.

3.3 Inference from search history

Although we could extend the above technique to the whole search we should recognise that information seeking is a dynamic activity; searches will change in focus, breadth and content over time. We should thus consider how components of this iteration are at explaining the current relevance assessments within the context of whole search. It is thus necessary to infer if the components are becoming more or less relevant over the search.

4. Conclusion

In our approach, RF is a process of *active selection* based on inference: inferring which documents, terms and characteristics of terms to use in RF. The inference mechanisms will select good components of explanations and weight them according to their value in explaining the relevance assessments. The resulting components can be assembled into different types of explanation according to the type of search.

Acknowledgements

The author gratefully acknowledges the support of the British Library LIC research grant number 991 for funding this work. More information is available at the project home page http://www.dcs.gla.ac.uk/ir/explanation/

References

[RB99] I. Ruthven and P. Borlund. *Information use and relevance assessments.* (in preparation). 1999

[RL99] I. Ruthven and M. Lalmas. S*elective relevance feedback using term characteristics.* CoLIS 3. Dubrovnik. 1999.