User-Chosen Phrases in Interactive Query Formulation for Information Retrieval

Alan F. Smeaton

School of Computer Applications
Dublin City University, Glasnevin, Dublin 9, IRELAND
asmeaton@compapp.dcu.ie

Fergus Kelleedy

Broadcom Éireann Research Ltd.
Kestrel House, Clanwilliam Place, Dublin 2, IRELAND
fk@broadcom.ie

Abstract

The impact of using phrases as content representation for documents and for queries has generally been accepted as a desirable feature in information retrieval systems because phrases are generally regarded as being more content-bearing than their constituent words. This has been borne by experiments in which the impact of phrases on retrieval performance has usually been found to be positive. However, most of the experimental results reported have derived phrases from documents and from queries in a fully automatic way. While this is acceptable for document indexing it is less acceptable for query formulation which is increasingly heading towards being an iterative process with users investing time in browsing the term space to choose appropriate search terms. In this paper we report a series of experiments in which two users, one experienced and the other a novice, formulate their queries by browsing the term space in advance of issuing a retrieval request. For these users we analyse the relative contributions and the impact of single words and multi-word phrases as search terms, on overall retrieval performance. Our results have implications for how choosing phrases as search terms should be presented to novice and to experienced searchers.

1. Introduction

Information Retrieval or resource discovery, is a “long-term, multi-threaded, iterative process with complex and dynamic requirements” [1]. This means that the information retrieval task itself has many dimensions, one of which is the fact that during query formulation users are not always clear about what they are looking for and even if they are they are not always sure about how to formulate that information need. In information retrieval systems one of the ways in which this can be accommodated to some degree is allowing the user to expand the initial query, possibly after some relevant documents have been found, in order to enrich the query specification.

The purpose of query expansion is to make the query resemble more closely the relevant documents and thus, hopefully, to retrieve those relevant documents. In IR research these days, query expansion could mean adding or deleting terms from the original query or even changing terms. This can be done using information from relevance feedback with relevant documents identified manually by the user [2] or by assuming the top-ranked documents from an initial ranking are relevant [3]. Other ways of discovering extra terms for a query involve using thesauri or lexical databases such as WordNet which have generally had mixed fortunes when used in IR [4,5]. Despite the apparent advantages offered by query expansion in terms of retrieval effectiveness, the uptake for the technique in operational information retrieval is quite limited and seems to be restricted to professional search systems, though there are some efforts at bringing this to a greater audience for searching the WWW [6].

Although query expansion goes some way towards helping users describe what they are looking for it is becoming clear to many in IR that an even more interactive query formulation processes is needed for the more explorative and undefined information needs users often have. This would exclude the typical cases of users hurriedly searching the WWW with 1- or 2-word queries to AltaVista or Excite but would include cases where users want to invest time in their searches, partly because they want to get value and partly because they are not sure what exactly it is they are looking for! The search paradigm advocated would be a 2-stage search process where the first part involves term navigation or query formulation, browsing the term space to choose terms and making precise the information need while the second part entails running the search and browsing through the results. Unfortunately, the real value of a 2-stage or even a multi-stage and iterative search process, in terms of

retrieval effectiveness has never really been explored in great depth in experimental information retrieval though the interactive track in TREC [7] has started to make inroads in this area. Experimental IR research has been dominated by evaluation of the single-shot retrieval mode, evaluating the effectiveness of matching the initial query against the document collection without iterating around the loop and evaluating the overall effectiveness of the process as is done in [8]. Some of this research has included evaluating the usefulness of phrases as indexing units for documents and for queries but the usefulness of phrases has always been measured as part of this one-shot retrieval model.

In this paper we will report a series of experiments in which we measure the contribution of phrases as indexing units in a retrieval strategy which involves the interactive formulation of a query. In the next section we will look at some other approaches to interactive query formulation and we follow that with a look at how phrase-based information retrieval affects performance. We then present details of our experimental environment and users, then our experimental results and finally some conclusions.

2. Interactive Query Formulation

As we know, formulating the query or search specification is a difficult task for users whose information needs are unclear, even to themselves but two-stage retrieval processes for IR are not new. The *Wide Area Information Servers* (wais) [9] had such an approach where the first stage involved the user choosing the actual set of globally distributed databases on which the search was to be performed. An idea reported in [10] is to use a stratified hypermedia browsing as developed by Bruza, as an information navigation task, to help a user form a query by browsing or navigating in the conceptual space to form several query fragments which can then be combined into one overall query. The domain of the search in this case is to search for facts but this is reminiscent of the work done at Siemens AG several years ago [11] where text documents were indexed based on an NLP analysis and the query navigation through the resulting structured lexicon, was the actual retrieval operation. Bruza himself has been working on the hyperindex concept for some time based on a 2-level model of an information space where on one plane there is the concept space of the terms and their inter-relationships, and these can “beam down” to the document space. In a web-searching implementation of this [12], the hyperindex can be viewed as a lattice of index expressions, dynamically constructed by retrieving the top 40 or 50 ranked web pages from a WWW search via a conventional web search engine. Recent work by Niwa *et al.* [13] has also presented an interactive guidance mechanism for document retrieval based on users browsing a visualised map of topics during query formulation.

What all these examples, and many more, show is that the concept of a 2-stage retrieval process is not new but it is not suited to all kinds of user searching. It is not always the case that users will gain some value from investing time and effort in this pre-search process because not all user searches require an exploration of the information need. It is only for cases where the information need needs clarification and interpretation that this process is worthwhile.

3. Phrase-Based Information Retrieval

A phrase can be defined as a concatenation of two or more words which must occur in text separated only by white space and do not range across paragraph or sentence bounds. As a unit of information, a phrase is more content-bearing than the sum of its constituent words and it is precisely because it is a richer representation of content than single words, phrase-based representations of document and query content can lead to improvements in retrieval effectiveness. As pointed out by Smeaton [14], the use of phrases in representing content can lead to difficulties in the matching of phrases during retrieval but for the most part, a phrase-based representation of content, coupled with a word-based representation, is seen as a worthwhile feature to be included in an IR system.

In the vast majority of the work reported in the IR literature on exploiting phrases for document retrieval, determining the phrases used to represent documents and queries is done as a static, automated process, with no user involvement. This is quite acceptable and indeed desirable for document indexing, but in most of the reported research on phrase-based retrieval the set of query phrases has come from automatically identifying phrases directly from the user’s input statement. Work has shown that even automatic phrase identification from a user’s topic statement makes positive contributions to improving retrieval effectiveness [15], however, using a users’ input string as a basis for query phrase determination is limiting and restrictive as automatic phrase identification depends on identifying phrases based on word adjacency or on linguistic relationships within the query. In many cases there are valid phrases which could be gainfully used in query representation but which would be conjunctions of words from different parts of the query input string or phrases which are composed of single words

not all of which occur in the original query.

There are currently two major ways in which a “phrase” can be identified. True linguistic phrases can only be identified linguistically and that requires some natural language processing to be performed on the text. As technology currently stands, this tends to be computationally expensive, domain-dependent and requires considerable auxiliary information such as lexicons, grammars, etc., in order to operate efficiently on large volumes of text [15]. This also requires that the linguistic analysis tool performs with reasonable accuracy and linguistic analysis is still an approximate science ! A much cheaper alternative to true phrases are “statistical” phrases which can be defined as collocations of 2 or 3 words which co-occur a significant number of times in a sizable portion of training text. Effectively this is based on the observation that phrases tend to occur in a corpus more often than not and if we pre-process a sample of text and count the number of times that co-occurrences actually occur, this gives us a phrase list. We can then use this list for subsequent identification of phrases in any text from the same domain as the sample.

An important series of experiments comparing the effectiveness of document retrieval based on using linguistic phrases, on using statistical phrases and on using no phrases at all, was performed by Fagan [15]. These experiments were repeated recently by Mitra *et al.* [16] also using SMART for retrieval but on a much larger data set of data set of about 250,000 documents and incorporating more recent and more effective term weighting functions. What makes the more recent results different to the runs that Fagan reported 10 years ago is that the effectiveness of the baseline performance in SMART has improved, the authors claiming that performance has almost doubled [16] and that this is due to new features being incorporated into retrieval like document length normalisation, new term weighting functions, etc. Mitra *et al.*’s argument is that statistical IR based on single word terms has simply got better in the last 10 years, possibly motivated by TREC. By re-examining statistical vs. syntactic phrases at this point in time, the landscape of how phrases can impact IR is completely different now to 10 years ago. The conclusions emanating from these runs are that adding statistical phrases to retrieval improves performance by 1%, whereas it used to be 7% and that contrary to Fagan’s findings, when phrases only (i.e. no single word terms) are used then syntactic phrases perform better than their statistical equivalent. Furthermore, using phrases does not significantly affect precision at the top ranks and phrases are good for some queries but not for others.

These results of Mitra *et al.* reflect a change in what we believed the impact of syntactic phrases really are, that they are indeed better than statistical phrases and single-word has got much better anyway, so the gap is negligible. However, the real impact of phrases on retrieval should not end there. The point about Mitra *et al.*’s work is that it assumes a completely automatic retrieval operation as it takes a TREC topic as input and automatically processes the query to derive query phrases. In operational IR, if a retrieval system takes a user’s input and makes some assumptions, like the derivation of any kind of phrases, then a user should know what is being assumed. Even this assumes a user inputs a query as a well-formed statement of information need like a TREC topic which is not what happens in the vast majority of operational IR searches. This work here really cries out to be complimented by a set of tests with real users who choose phrases from a phrase list, in response to a query input. This should then rightly be compared against having a user explore the concept space of words only and let the user choose related words, given the same amount of time in query formulation. However, the selection of phrases for expansion above is easier than the expansion by words as the full topic description gives a handle on a set of phrases to browse from the phrasal lexicon, i.e., all those phrases containing at least one of the topic words.

As an example of such an approach, [17] describes a retrieval system which presents the result of a query as a ranked list of contexts or concepts or parts of the document space, which best match the query. The user then chooses one of these contexts and from within this context manipulates the original query by perhaps adding new terms (words or queries) before submitting an extended query to the entire database. So instead of a query yielding a document subset to peruse among, it yields an information and term space to browse among before retrieving in a much more informed manner via an expanded query against the whole corpus. What is interesting about this is that the query manipulation within a context includes the ability to pull down phrases from a phrase list to add to the original query and there is great justification in the paper for such a process.

In the present paper we set out to explore the impact of statistically identified phrases in an interactive query formulation environment. This is distinct from other reported work [15,16,18] including our own [19] where phrases are identified from user queries automatically. Marchionini has argued that users seek the path of least cognitive resistance and prefer recognition tasks to recall tasks [20]. This explains why most users search by initially entering short, imprecise queries and are happy to browse among many non-relevant documents retrieved instead of adopting the search strategy which has been experimentally proven to work, that if we spend time on

query formulation then we will benefit. Though we know this, it is harder for us to do so we don't do it and we take the easy route. If a retrieval system forces user dialogue at query formulation by offering candidate phrases from a phrase list then this is cognitively less stressful to a user than having to formulate phrases, recognition being easier than recall. Thus retrieval based on an initial short query, followed by an invitation to browse/choose from a generated phrase list, should be inherently easier for a user and here we set out to find out how this affects overall retrieval effectiveness.

4. Experimental Environment: Documents, Topics and Users

The set of experiments we report in this paper involve our use of a test collection of documents, queries and relevance assessments. For our experiments we used 74,520 articles from the Wall Street Journal which we indexed by removing stopwords and stemming the remaining words, using the word stems as part of the representation for documents. We then applied a simple phrase recognition algorithm to the documents, generating 2-word phrases from adjacent occurrences of non-stopwords within the same sentence, and generating 3-word phrases from adjacent occurrences of non-stopwords or from occurrences of non-stopwords separated by an occurrence of one stopword, within the same sentence. Thus the string “The Department of Defense overturned appeals from ...” would yield the phrases

<i>department_of_defense</i> ¹	<i>defense_overturn_appeals</i>
<i>defense_overturn</i>	<i>overturn_appeal</i>

An important part of the phrase indexing routine we used is not to index text by all phrases occurring within it but only by those phrases that are meaningful. In the case above, the phrase *defense_overturn* is not really a content-bearing part of the text, yet it is a valid phrase. To help improve the quality of phrases used in indexing we decided to index documents only by valid phrases from the document which had occurred 25 times or more from a pre-processed sample of 260 Mbytes of text. This “lexicon” of 91,964 valid phrases would hopefully not contain the entry *defense_overturn* (in fact it does not) but would contain entries for the other phrases because of the likelihood of them occurring in text anyway. Further details of our phrase recognition process can be found in [19]. Once documents have been indexed by word stems and by phrases they were input into our own information retrieval system, developed in previous work [21] and retrieval was based on assigning $tf*IDF$ weights to query terms and scoring documents as the sum of the term weights of query terms (words and phrases) which index them.

For queries, we used 50 topics from TREC-5 and we allowed 2 different users to read and understand the topic and then to input into our system some initial text, a list of words or a natural language input if they wished. Each of the non-stopword query terms was used as a query term. The users were then presented with access to a sorted list of valid phrases from the phrasal lexicon and a navigation tool through this list to help them quickly locate and choose to use, phrases as part of the query representation. Each user was prompted with the list of phrases which contained two or three of their original query terms, as well as phrases from the lexicon containing only one of their indicated query terms and users were free to add valid phrases to the query as they pleased though not phrases containing none of their original query words as constituents. When the query formulation task was complete, the list of word stems and valid phrases was used as a query representation in a term-weighted retrieval strategy described in [21].

In our experiments we used two different users who each performed the query generation task independently of each other. One user (UserA) could be termed *internet-savvy* with a good deal of experience in information retrieval techniques while the second user (UserB) was a complete novice to searching using an IR search engine and had never completed even a search of the WWW. For each user's representation of each query, we ran the query against the corpus of documents and using the set of relevance judgments from TREC-5 we were able to compute precision-recall figures for each user's query formulation. For the purposes of comparison of retrieval across different runs, we computed precision at rank positions 5, 10, 20 and 30, as well as average precision across all retrieved documents and the complete precision-recall curves, using standard TREC techniques. The results of our runs are reported in the next section.

¹ For clarity of explanation, we show here the phrases represented as concatenations of words whereas in fact we use concatenation of word stems.

5. Experimental Results

The first set of retrieval runs we did was to take the queries generated by UserA and by UserB and run them against the corpus, keeping the user's original words and queries in the query. We then re-ran the experiments using only the words from the query and then using only the phrases from the query to measure the relative contributions of each. The performances of these variations of queries are shown in Figure 1 and Figure 2 shows the complete precision-recall curves for all runs. These results do not appear too performant at a qualitative level because we have not incorporated any of the “smarts” known to improve IR performance such as proper name recognisers or document length normalisation.

Query type	UserA			UserB		
	Words & Phrases	Words	Phrases	Words & Phrases	Words	Phrases
P @ 5	0.1880	0.1960	0.1280	0.1440	0.2040	0.0640
P @ 10	0.1400	0.1600	0.1040	0.1160	0.1560	0.0460
P @ 20	0.1230	0.1200	0.0910	0.0920	0.1140	0.0370
P @ 30	0.1040	0.0980	0.0767	0.0767	0.0933	0.0307
Avg. P	0.1770	0.1574	0.0758	0.1522	0.1489	0.0275

Figure 1: Performance figures for experimental runs

Figure 2: Precision-Recall curves for experimental runs

There are many interesting points to be drawn from these results, the first of which is the fact that words alone perform significantly better than phrases alone when used as query terms, for both users. The second observation is that userA performs better at query formulation than userB, marginally better for words and significantly better

for phrases. In fact the phrases chosen by userB (the naïve user) are quite poor and when combined with userB's words actually bring down retrieval effectiveness at high precision. The most significant and surprising result however is that for userB, words alone perform better than words and phrases at the high precision end, though not for precision averaged over the whole ranking while for the experienced userA, words and phrases were better across the whole ranking than words or phrases, but at the very high precision (at 5, 10 documents), words alone were better than words and phrases. This result is a bit unexpected. To examine why this might be so we computed the overlap in words and in phrases chosen by the two users for the set of queries. These are shown in Figures 3 and 4.

Figure 3: Overlap in Words Chosen by UserA and UserB

Figure 3 shows a considerable overlap in the single word query terms chosen by the two users which is not surprising. By virtue of having slightly more unique words per query (8.7 words on average), userA's queries contain marginally more single word terms than those from userB (7.7 average) but with 7.3 words in common. The amount of overlap in phrases chosen by the two users as shown in Figure 4 is, however, very small. UserA had an average of 18.5 phrases per query while userB had a comparable average of 13.9 phrases per query yet there were only an average of 2.8 of these common across the two sets.

These figures do not explain the unexpected results we have obtained so we investigated further the impact of adding phrases to the query by computing the length of the posting list entries for the single word and phrase terms used by userA and by userB. These show that the average postings length for words in the queries were 3976 for userA and 3866 for userB while the average postings length for the phrases added by the users were only

43.6 for userA and 44.5 for userB. These figures point to the much diluted impact that adding phrases actually has on the set of documents being scored, remembering that there are 74,520 documents in the collection. This means that when a phrase is added to a query it impacts on only a small fraction of the document scores because phrases naturally occur much less frequently than single word terms.

Figure 4: Overlap in Phrases Chosen by UserA and UserB

6. Conclusions

Our experiments have thrown up a surprising result. As expected, for the experienced user (userA), benefit was obtained by allowing that user to interactively browse the phrasal lexicon and to select phrases to add to the query. For the naïve user (userB) this was not found to be the case and adding phrases in fact hindered performance over just using words alone. Even though these experiments were carried out on only 2 users and the present work should be validated on a larger population, we believe that the results we have found are important enough to be reported. Because the postings list entries for phrases tend to be much shorter than those for single word terms, i.e. phrases occur much less frequently than single words, the impact on overall retrieval performance across the whole ranking is reduced.

Much work in information retrieval in recent times has highlighted the fact that two or more independent formulations of the same query, or two independent rankings of documents using different retrieval strategies, when combined, tend to yield an improvement in retrieval effectiveness over any of the inputs. This was the case when the single-word only and the phrase-only formulations of the queries were combined for userA into one

ranking where words and phrases performed better than words or than phrases alone. However, for userB, when generating two different document rankings from two different formulations of the query (one using words only, the other using phrases only) and then combining these into one ranking by retrieving based on words and phrases, the combined result was significantly worse than retrieval based on words alone. This suggests that the phrases chosen by userB were simply badly chosen and further suggests that a two-stage process for retrieval where users spend stage one browsing the term space adding word or phrasal index terms to their query, followed by a retrieval operation and a browse among the document ranking, may only be appropriate for those users who have more than basic training. For our userB we deliberately chose a non-computing undergraduate student, majoring in economics, with no experience in information retrieval or searching at any level.

The fact that phrase-based retrieval worked well in our case only for the experienced user has implications for how phrase-based retrieval is incorporated into operational IR. The effectiveness results we have obtained mirror those found by others when they use **automatic** determination of search phrases. In the most recent of these results, Mitra *et al.* [15] have found that using phrases does not significantly affect precision at the top rank positions and that phrases are good for some queries but not for others, exactly as we have found.

To integrate phrases more effectively into the IR process may require a different approach to that we have taken here. Phrases and single word terms have different frequency distributions and this should really be taken into account in applying 2 different term weighting functions as part of retrieval instead of bundling the phrase and single word representations into one as we have done. Logically, the term weighting does not sit comfortably with treating words and phrases as equivalent terms and more work is required in the basic IR strategy which uses phrases. This could, for example, downweight in some way the impact of single word terms X and Y occurring in a document if that document also has the phrase XY.

The fact that our search phrases were chosen manually as opposed to automatically and we obtain the results we have, is disappointing and requires further investigation. For the present, how we could use phrase searching in the most popular of information retrieval applications for the untrained user, searching the WWW also deserves careful consideration. Our work has been concerned with using phrases in non-professional searching scenarios, eg searching the WWW and for novice searchers neither the automatic determination of query phrases from natural language queries nor our approach of a 2-stage phrase selection process, seem to be appropriate solutions for all users. If phrase based information retrieval is to become widespread then this needs to be addressed.

Acknowledgements:

The authors would like to acknowledge the comments of Stephen Robertson and Ross Wilkinson as well as the two anonymous referees whose comments have helped to improve this paper.

References

1. Lagoze, C. From Static to Dynamic Surrogates: Resource Discovery in the Digital Age. D-Lib Magazine, June 1997, <http://www.dlib.org/dlib/june97/06lagoze.html>
2. Smeaton AF and van Rijsbergen CJ. The Retrieval Effects of Query Expansion on a Feedback Document Retrieval System. The Computer Journal, 26(3):239-246, 1983.
3. Buckley C., Singhal A., Mitra M. and Salton G. New Retrieval Approaches Using SMART: TREC-4. In Proceedings of TREC-4. DK Harman (Ed.). NIST Special Publication 500-236, 1996.
4. Voorhees, EM. Using WordNet to Disambiguate Word Senses for text Retrieval. In: Proceedings of the 16th ACM-SIGIR Conference, Pittsburgh, 1993:171-180
5. Voorhees, EM. Using WordNet for Text Retrieval. In: C. Fellbaum (Ed.) WordNet: An Electronic Lexical Database and some of its Applications. MIT Press, 1996.
6. Crimmins F and Smeaton AF. Relevance Feedback and Query Expansion for Searching the Web: A Model for Searching a Digital Library. In: Proceedings of First European Conference on Research and Advanced Technology for Digital Libraries, ECDL'97, Springer Notes in Computer Science #1324:99-112, 1997.

7. Over P. TREC-5 Interactive Track Report. In: Proceedings of the Fifth Text Retrieval Conference (TREC-5), EM Voorhees and DK Harman (Eds.):29-56, 1997.
8. Magennis, M. ??? In: Proceedings of the ACM SIGIR Conference, Philadelphia, July 1997:????
9. Pfeifer U. Searching Structured Documents with the Enhanced Retrieval functionality of free-WAIS sf and Sfgate. U Dortmund Lehrstuhl Informatik VI, D-4221, <http://www.fhg.de/www/www95/paper/47/fwsf/fwsf.html>.
10. ter Hofstede AHM Proper HA and van der Weide ThP. Query Formulation as an Information Retrieval Problem. The Computer Journal, 39(4):255-274, 1996.
11. Thurmair G. A Common Architecture for Different Text Processing Techniques in an Information Retrieval Environment. Proc ACM Conference on R&D in IR, Pisa, Sept 1986:138-143.
12. Bruza PD and Dennis S. Query Reformulation on the Internet: Empirical Data and the Hyperindex Search Engine. In: Proceedings of RIAO97, Computer-Assisted Information Searching on Internet, Montreal, June 1997.
13. Niwa, T *et al.* Topic Graph Generation for Query Navigation: Use of Frequency Classes for Topic Extraction. In: Proceedings of NLPRS'97, Natural Language Processing Pacific Rim Symposium, Phuket, Thailand, 1997.
14. Smeaton, AF. Progress in the Application of Natural Language Processing to Information Retrieval Tasks. The Computer Journal, 35(3):268-278, 1992.
15. Fagan JL. Automatic Phrase Indexing For Document Retrieval: An Examination Of Syntactic And Non-Syntactic Methods. In Proceedings of the Tenth ACM SIGIR Conference on Research and Development in Information Retrieval, New Orleans, June 1987:91-108.
16. Mitra, M. *et al.* An Analysis of Statistical And Syntactic Phrases In: Proceedings of RIAO97, Computer-Assisted Information Searching on the Internet, Montreal, Canada, June 1997:200-214.
17. Anick PG and Vaithyanathan S. Exploiting Clustering and Phrases for Context-Based Information Retrieval in: Proceedings of ACM SIGIR Conference, Philadelphia, July 1997:314-
18. Smeaton, AF. Using Parsing of Natural Language as Part of Document Retrieval, PhD thesis, University College Dublin, 1987.
19. Kelledy, F and Smeaton, AF. Automatic Phrase Recognition and Extraction from Text. In Information Retrieval Research: Proceedings of the 19th Annual BCS-IRSG Colloquium on IR Research, Aberdeen, Springer Electronic Workshops in Computing, 1997.
20. Marchionini, G. Interfaces for End-User Information Seeking. Journal of the ASIS 43(2):156-163, 1992.
21. Kelledy F. and. Smeaton AF. TREC-5 Experiments at Dublin City University: Query Space Reduction, Spanish Stemming and Character Shape Coding. In: Proceedings of TREC-5, NIST Special Publication 500-238:197-207, 1997.