Structure of morphologically expanded queries:
A genetic algorithm approach

Lourdes Araujo [a,*], Hugo Zaragoza [b], Jose R. Pérez-Agüera [c], Joaquín Pérez-Iglesias [a]

[a] Universidad Nacional de Educación a Distancia (UNED), Madrid, Spain
[b] Yahoo! Research, Barcelona, Spain
[c] Universidad Complutense de Madrid, Spain

ARTICLE INFO

ABSTRACT

In this paper we deal with two issues. First, we discuss the negative effects of term correlation in query expansion algorithms, and we propose a novel and simple method (query clauses) to represent expanded queries which may alleviate some of these negative effects. Second, we discuss a method to optimize local query-expansion methods using genetic algorithms, and we apply this method to improve stemming. We evaluate this method with the novel query representation method and show very significant improvements for the problem of stemming optimization.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

User queries to retrieve information from the web or a document collection are very often too short or ambiguous. Query reformulation is a useful technique to improve the retrieval results of those queries. There is an underlying common background in many of the works done in query reformulation, namely the appropriate selection of a subset of search terms among a list of candidate terms. However, we cannot evaluate a priori the quality of a subset with respect to another one: this depends on the (unknown) relevance of the documents in the collection and the number of possible subsets grows exponentially with the size of the candidate set. For these reasons, standard optimization techniques cannot be applied to this problem. Instead, we must resort to heuristic optimization algorithms, such as genetic algorithms, that sample the space of possible subsets and predicts their quality in some unsupervised manner.

Before considering any query reformulation process it is very important to take into account that modern information retrieval ranking functions apply the *term independence* assumption. This assumption takes on many forms, but loosely it implies that the *effect of each query term on document relevance can be evaluated independently of the other query terms*. This has the effect of rendering all queries *flat*, without structure.

However, there are many cases in which queries have some known *linguistic structure*, such as degree of synonymy between terms, term co-occurrence or correlation information with respect to the query or to specific query terms. This is typical of queries constructed by a query-expansion method, of stemming or normalizing terms, of taking into account multi-terms or phrases, etc. Surprisingly, almost all ranking functions (and experiments) ignore this structural information: after expansion, selection and re-weighting of terms, a flat query (a set of weighted terms) is given to the ranking function which *assumes terms are independent* and scores documents accordingly.

---

* Corresponding author.
E-mail addresses: lurdes@lsi.uned.es (L. Araujo), hugoz@yahoo-inc.com (H. Zaragoza), jose.aguera@fdi.ucm.es (J.R. Pérez-Agüera), joaquin.perez@lsi.uned.es (J. Pérez-Iglesias).

Our goal is to propose a new method that solves the problems related to the selection of suitable terms for query reformulation and term independence assumption. The morphological query structure of the query has been chosen to show how our approach is capable to improve state-of-art approaches like Porter's stemmer.

In Section 2 we propose a novel way to represent expanded queries that encodes information about the term correlation using clauses like set of related terms. The proposed representations greatly increase the expressiveness power of queries, but at the expense of introducing parameters (weights) which may be hard to set. Section 3 shows one experiment where our clauses representation model is adapted to the problem of morphological query expansion. In Section 4 we apply it to the problem of optimizing the expansion of a term with respect to its stem. We show that we can significantly improve the performance of Porter stemming by adapting the expansion to every query. Section 5 is devoted to describe some related works and Section 6 draws the main conclusions and describe the future lines of work.

2. Ranking independent clauses of dependent terms

One of the reasons of the high performance of modern ad hoc retrieval systems is their use of document term frequency. It is well known [22] that (i) probability of relevance of a document increases as the term frequency of a query term increases, and (ii) this increase is non-linear. For this reason most modern ranking functions use an increasing saturating function to weight document terms that are in the query. An example of this is the term saturating function used as part as BM25 [22]:

$$w(d,t) := \frac{tf(d,t)}{tf(d,t) + K1} \tag{1}$$

where $tf(d,t)$ is the term frequency of term t in document d, and $K1$ is a constant. Similar non-linear term frequency functions are found in most IR ranking models such modern variants of the vector space model, the language model, divergence from randomness models, etc. Besides, all these ranking functions assume that *the relevance information of different query terms is independent* and therefore the relevance information gained by seeing query terms can be computed separately and added linearly (or log-linearly), for example, in BM25:

$$score(d) := \sum_{t \in q} w(d,t) \cdot idf(t) \tag{2}$$

This independence assumption is usually reasonable for short queries (i.e. "*Italian restaurant in Barcelona*"), since users use each term to represent a different aspect of the query. However, such assumption breaks down for queries that are sufficiently complex to contain terms with sufficiently close meaning. Consider for example the query "*Italian restaurant cafeteria bistro Barcelona*". Having seen the term *restaurant* twice in a document, which term is more informative: *Barcelona* or *cafeteria*? Loosely speaking, if a group of terms carries the same meaning, the amount of relevance information gained by their presence should diminish as we see other terms in this group, very much like in Eq. (1) does for term frequency, and unlike (2).

This situation arises very often in modern IR tasks and systems, in particular in the following areas:

- morphological expansion (e.g. stemming, spelling, abbreviations, capitalization),
- extracting multi-terms from the query,
- query term expansion (e.g. user feedback, co-occurrence based expansion),
- lexical semantic expansion (e.g. using WordNet),
- using taxonomies and ontologies to improve search,
- user modeling, personalization,
- query disambiguation (where terms are added to clarify the correct semantic context),
- finding similar documents (where the query is an entire document),
- document classification (where the query is a set of documents),
- structured queries (such as TREC structured topics).

We propose to consider two *levels* of representation: *terms* and *term clauses*. Clauses are *sets of weighted terms* that are intended to represent a particular aspect of the query. The weights represent their relative importance within the clause (in particular, the strength of the dependence with relevance). Thus, a query can be thought of as a bag of bags of (weighted) terms:

$$c := \{(t_0, w_0), (t_1, w_1), \dots, (t_{|c|}, w_{|c|})\}$$
$$q := \{c_1, c_2, \dots, c_{|q|}\}$$

Boolean retrieval models and the Inquery [3] retrieval model have used query representations even more general than this. Here we restrict ourselves to this representation with two levels to give clear semantics to each level: term and clause. We are going to consider terms *within* a clause as if they were greatly dependent with respect to relevance; in fact we will consider them as if they were virtually the same term. Second, we consider terms *across* clauses as being independent with respect relevance, as is usually done across terms.

Conceptually, what we propose is a projection from the space of terms to the space of clauses. Formally, we represent a document as the vector $d = (tf_1, \ldots tf_i, \ldots, tf_V)$ where V if the size of the vocabulary. We represent a query having n clauses as a $n \times V$ matrix of weights: $C = (c_{ij})$ where c_{ij} is the weight of jth term in ith clause. The projected document is then $d|_C := d \times C^\mathsf{T}$.

Consider this example. Imagine that we are given a corpus with four terms A–D:

Doc	A	B	C	D
d_1	2	1	0	1
d_2	0	1	1	1
d_3	1	2	0	2

Now consider the query with two clauses:

$$q := \{\{(A, 1.0), (B, 0.7)\}, \{(C, 1.0)\}\}$$

This query can also be represented by the matrix:

$$C := \begin{bmatrix} 1 & .7 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

In the projection $d|_C := d \times C^\mathsf{T}$ the original query terms are removed from the collection and replaced with new pseudo-terms representing the clauses; other terms are removed because all terms are included in the clauses. The term frequency of the clause pseudo-term will be equal to the weighted sum of term frequencies of the terms in the clause. In our example:

Transformed Doc	Clause 1	Clause 2
$d_1|_q$	2.7	0
$d_2|_q$	0.7	1
$d_3|_q$	2.4	0

This projection will be different for every clause, so it cannot be done as a pre-processing step: it must be done online. In practice one can carry out this transformation very efficiently since the projection is performed only on the few terms contained in the query, and document vectors are very sparse. Most important, the information needed for this projection is contained in the postings of the query terms. This means that one can compute the needed term statistics on the fly after postings are returned from the inverted index. This will incur the cost of a few extra flops per document score, but without any extra disk or memory access.

The length of the document is not modified by the projection, nor the average document length. The *clause term frequencies* (*ctf*) and *clause collection frequencies* (*ccf*) can be computed as:

$$ctf(d, c) := \sum_{(t,w) \in c} w \cdot tf(d, t)$$

$$ccf(d, c) := \sum_{(t,w) \in c} w \cdot \sum_{d} tf(d, t)$$

being d the document, c the considered clause, and w a weight, for example it can be the re-weighting value assigned to the terms in query expansion tasks. In our case, this weight always takes the value 1.

From these measures, we can define the probability of a clause in a document, $p_{\mathrm{ML}}(c|d)$, and within the collection, $p_{\mathrm{ML}}(c|Col)$, in the following manner:

$$p_{\mathrm{ML}}(c|d) := \frac{ctf(d, c)}{ctf(d, c) + \sum_{t \notin c} tf(d, t)}$$

$$p_{\mathrm{ML}}(c|Col) := \frac{ccf(d, c)}{ccf(d, c) + \sum_{d, t \notin c} tf(d, t)}$$

The most problematic statistic is the *inverse clause frequency* (*icf*), since this is not clearly defined in the weighted case. One possible choice is the number of postings containing *at least one term in the clause*; we refer this as $icf_{\mathrm{OR}}(c)$, and we note that it can be computed directly from the size of the clause result set (documents with non-zero *ctf*). However, this number may be unfairly large for clauses with lowly weighted common terms. Furthermore, in some settings this number may not even be available (for example if we only score the query term AND set or if we drop from the computation documents unlikely to be highly scored). Another possibility is to use the *expected idf* for a clause term in a document:

$$icf_E(d,c) = \frac{1}{ctf(d,c)} \sum_{(t,w) \in c} w \cdot tf(d,t) \cdot idf(t) \qquad (3)$$

In our empirical evaluation we found this is better than using the *min* or the *max clause idf*, and better than using the *mean idf*.

With these statistics at hand we can compute the relevance score of a document with respect to a query with clauses for a number of retrieval systems [2]; we display several in Table 1.

3. Query clause experiments

We have performed experiments to demonstrate the dangers of the term independence assumption for queries with strongly correlated terms, and to test the proposed *query-clauses* ranking idea applied to the stemming problem. Evaluation has been carried out on the Spanish EFE94 corpus which is part of the CLEF collection [18] (approximately 215 K documents of 330 average word length and 352 K unique index terms) and the 2001 Spanish topic set, with 49 topics of which we only used the title (of 3.3 average word length). All runs employ the standard (Eq. (1)) and the query clause version of BM25 Table 1.

3.1. Stemming

One can view stemming as a form of global query expansion: we expand a term in the query with every term in the morphological thesaurus sharing the same stem. However, doing this directly on the query hurts performance – mean average precision – greatly (Table 2, rows 1 and 2). One may think that this loss of performance is due to the noise introduced by the stemming algorithm, but this is not the case: if we replace terms by their stemmed version at indexing time, performance will most often increase and rarely decrease (Table 2, row 3). In our opinion, this is due to the strong violation of the term independence hypothesis produced by adding so many strongly correlated terms at query time.

A natural way to expand a query by stemming is to construct a set of sets of terms, or a set of *clauses*, where each clause represents all the forms of a stem, possibly weighted (since we may want to weight more strongly the original term typed by the user). The resulting query is a set of sets of clauses which can be ranked with our proposed method. Its performance, using as *idf* the clause's (icf_{OR}) is exactly equivalent to stemming the collection since in both cases term frequencies of stems are collapsed (this is also seen empirically in Table 2, last row).

3.2. Our genetic algorithm

Because we need to perform the selection of a particular set of terms among a huge amount of possible combinations of candidate query terms, the computational complexity of exhaustive search methods is non-viable and we have resorted to a heuristic method such as a genetic algorithm.

Genetic algorithms [12] have been shown to be practical optimization methods in very different areas [17]. Evolutionary algorithms mimic the principles of natural evolution: heredity and survival of the most fit individuals.

A genetic algorithm maintains a population of potential solutions, and is provided with some selection process based on fitness of individuals. The population is renewed by replacing individuals with those obtained by applying "genetic" operators to selected individuals. The usual "genetic" operators are *crossover* and *mutation*. Crossover obtains new individuals by mixing, in some problem dependent way, two individuals, called parents. Mutation gives a new individual by performing some kind of change on an individual. The production of new generations continues until resources are exhausted or until some individual in the population is fit enough. Fig. 1 shows the structure of a genetic algorithm. The algorithm works with a collection of individuals $\{x_i, \ldots, x_n\}$, called population. Each individual represents a potential solution to the problem considered, implemented as some data structure, which depends on the problem. The evaluation of each solution gives a measure of its *fitness*. At a new generation step, a new population is formed by selecting the more fit individuals. Some members of the new population suffer transformations as a consequence of applying *genetic operator* to form new solutions. After a number of generations, the program is expected to converge, and it is hoped that then, the best individual represents a solution close to the optimum.

Chromosomes of our GA are fix-length binary strings where each position corresponds to a candidate query term. A position with value one indicates that the corresponding term is present in the query. Fig. 2 shows an example. Because some

Table 1
Implementing query clauses in several standard ranking models.

Model	Weighting
BM25	$\frac{ctf}{ctf+K} \cdot icf$
VSM	$\frac{ctf \cdot icf}{\|d_{iq}\|}$
DFR (PL2)	$\frac{1}{ctf+1}\left(ctf \cdot \log_2 \frac{ctf}{\lambda} + (\lambda - ctf) \cdot \log_2 e + 0.5 \cdot \log_2(2\pi \cdot ctf)\right)$
LM (KL)	$p_{\text{smoothed}}(c\|q) \log(p_{\text{smoothed}}(c\|d))$

Table 2
Stemming performance. Avg. Prec. stands for mean average precision and Prec10 for the precision of the first 10 documents retrieved.

Method	Avg. Prec.	Prec10
No Stemming	.37	.47
Stem Expansion (Standard)	.20	.28
Stemming	**.43**	**.52**
Stem Expansion (Clauses)	**.43**	**.52**

```
evolution program
begin
    generation = 0
    P = initialize_population
    F = evaluation(P)
    while not required_fitness(F) and
          not termination_condition do
    begin
        generation = generation + 1
        I = individuals_selection(P, F) %for genetic operators
        P = new_generation(P, I)
        F = evaluation(P)
    end
end
```

Fig. 1. Structure of a genetic algorithm.

preliminary experiments performed have shown that, in most cases, the elimination of the original query terms degrades the retrieval performance, we force to maintain them among the selected terms of every individual. The set of candidate terms is composed of the original query terms, along with related terms provided by the applied morphological thesaurus. Each term of the original query is grouped with the expanded terms related to it, in the called *term_set* [20]. The weights assigned to the documents retrieved with each *term_set* are used to sort the total set of retrieved documents in the fitness function.

The applied selection mechanism has been roulette wheel. In roulette wheel selection, the chances of an individual to be chosen for reproduction are proportional to its fitness. We apply the one-point crossover operator and random mutation [11,12,17]. In one-point crossover a single crossover point is chosen on both parents strings. The parts of the parent strings divided by the crossover point are swapped to generate two children, containing a part of each parent. The random mutation operator simply flips the value of a randomly chosen bit (0 goes to 1 and 1 goes to 0). We also apply elitism, the technique of retaining in the population the best individuals found so far. The fitness function used is some measure of the degree of similarity between a document belonging to the system and the submitted query. We will discuss this further in the different experiments.

There have been several works applying GAs [9] to different aspects of information retrieval, and also to the stemming problem [10]. Proposals devoted to the query expansion problem with GAs can be classified into relevance feedback techniques and Inductive Query by Example (IQBE) algorithms. In systems based on relevance feedback [26] the user gives feedback on the relevance of documents retrieved by his original query. IQBE [6] is a process in which the user does not provide a query, but document examples and the algorithms induce the key concepts in order to find other relevant documents.

Candidate query terms:

ballena	ballenas	reserva	reservada	reservaba	reservan	reservar	reserven	...

Population:

I1:

1	1	1	0	0	0	0	1	...

I2:

1	1	0	0	1	0	0	0	...

I3:

1	1	1	0	1	1	0	1	...

I3: ...

Fig. 2. Example of some individuals of the GA population of the Spanish query *reserva de ballenas* (whale reserve).

Table 3

Global precision results for the whole set of tested queries. Each individual datum has been computed as the average over five different GA runs. Prec. stands for precision (all documents), Prec10 stands for the precision of the results for the first 10 documents retrieved. Last column is the rate of precision (all documents) improvement.

	Prec.	Prec10	Improvement (%)
Baseline	0.3567	0.4150	–
Porter Stem.	0.4072	0.45	+12.40
Genetic Stem.	0.4682	0.54	+23.81

Mainly two information retrieval problems have been tackled with GAs: assigning weights to the query terms [21,28,23,14,13], and selecting query terms. Let us consider a number of proposals in the latter case, the one on which we focus our work. Chen et al. [6] apply a GA as an IQBE technique, i.e. to select the query terms from a set of relevant documents provided by the user. In this work, the authors propose an individual representation that has also been used in later works: chromosomes are binary vectors of fixed size in which each position is associated with one of the candidate terms. In [15], the authors propose a genetic programming (GP) algorithm to learn boolean queries encoded as trees whose operators are AND, OR and NOT. This work was later extended in [7] by incorporating multiobjective optimization techniques to the problem. Fernández-Villacañas and Shackleton [16] compared two evolutionary IQBE techniques for boolean query learning, one based on GP and the other on a classic binary GA, which obtained the best results. Kraft et al. [19] propose the use of GP to learn fuzzy queries. The queries are encoded as expression trees with boolean operators as inner nodes, and query terms with weights as terminal nodes. Cordón et al. [8] extend Kraft's proposal by applying a hybrid simulated annealing-genetic programming scheme, what allows them to use new operators to adjust the term weights. Tamine et al. [25] use knowledge-based operators instead of the classical blind operators, as well as niching techniques.

A common factor of the above mentioned works is that they relay on some kind of information provided by the user. In some cases, the user has to provide a set of documents that are used for the inductive learning of terms. In other cases, the user provides relevance judgements on the retrieved documents, that are use to compute the fitness. On the contrary, we have studied the GA which is applied without user supervision.

3.3. Testing the viability of the approach

Our first experiments have been aimed to test how much the GA can improve the results of a search. To this purpose, we take the user relevance judgements as the best fitness function that we can use. Since for the CLEF collection used in the experiments we have the user relevance judgements, we have used them to guide the selection process. Specifically, we have used as fitness function the standard precision measure defined as the fraction of retrieved documents (the set A) which are relevant (R_a).

$$\text{Precision} = \frac{|Ra|}{|A|}$$

Table 3 shows the precision of the best individual obtained by the GA. We have compared our system performance with the results of the original user query (*Baseline*) and with the results obtained expanding with the stems provided by the Porter stemming (*Porter Stemming*). The latter works by substituting the original query terms by their stems and performs the search in the document collection indexed by stems. We can observe that our system achieved an important improvement of the performance, greater than the one achieved with other stemming methods traditionally used in query expansion, such as Porter. Fig. 3 shows the evolution of two queries of the test set. We can observe that both of them reach convergence very quickly.

4. Experimental results

The system has been implemented in Java, using the JGAP library,[1] that provides a generic implementation of a genetic algorithm, on a Pentium 4 processor.

We have carried out experiments to evaluate the fitness functions considered. We have investigated the best range of the GA parameters. Finally, we provide global measures for the whole set of queries that we have considered.

4.1. Selecting the fitness function

In the beginning, we would like to use Average precision as the fitness function. However, this is not known at query time. Instead, it has been suggested in previous work to use the document scores as the fitness [14]. While this may not be intuitive, it turns out that variations of these scores after expansion are correlated with relevance [1]. One intuitive explanation
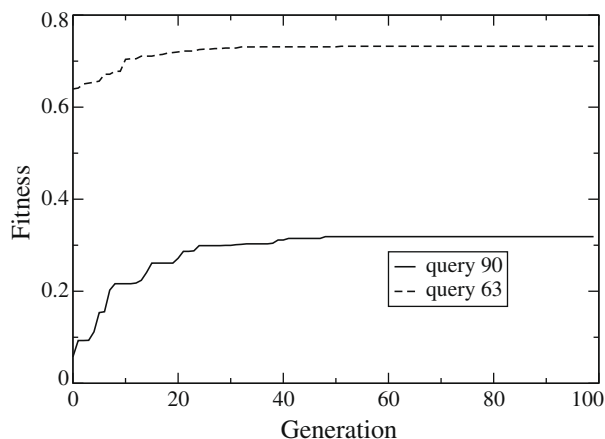
---

[1] http://jgap.sourceforge.net/.

Fig. 3. Fitness evolution for two queries of the test set. The GA parameters have been a population size of 100 individuals, a crossover rate of 25% and a mutation rate of 1%.

would be that adding an unrelated term to a query will not bring in new document with high scores, since it is unlikely that it will retrieve new documents; on the other hand adding a term that is strongly related to the query will bring new documents that also contain the rest of the terms of the query and therefore it will obtain high scores.

We have considered three alternative fitness functions, $\sqrt{\cos\theta}$, $\cos\theta$ and $\cos^2\theta$. To select the fitness function to be used in the remaining experiments, we have studied the fitness evolution for different queries of our test set. Fig. 4 compares the fitness evolution for the query which reaches the greatest improvement (*best_query*). The three functions converge to different numerical values that correspond to the same precision value (.68). We can observe that the square-root cosine function if the first one to converge. A similar behavior is observed in other queries. Accordingly, the square-root cosine has been the fitness function used in the remaining experiments.

4.2. Tuning the GA parameters

The next step has been tuning the parameters of the GA. Fig. 5 shows the fitness evolution using different crossover (on the left) and mutation (on the right) rates for the best query. Results show that we can reach a quickly convergence with values of the crossover rate around 25%. Mutation rates values around 1% are enough to produce a quick convergence.

Fig. 6 show the fitness evolution for the best (on the left) and the worst (on the right) queries, with different population sizes. The plots indicate that small population sizes, such as one of 100 individuals, are enough to reach convergence very quickly.

4.3. Overall performance

Table 4 (Stem Expansion (Clauses)) shows the results obtained using stemming as query expansion but building a clause for every term. As expected, the results are exactly those obtained in traditionally stemming by collapsing terms to their stems.

In order to show the improvement of our approach, we have compared the genetic algorithm performance with the results of the original user query (*Baseline*) and with the results obtained expanding with the stems provided by the Porter
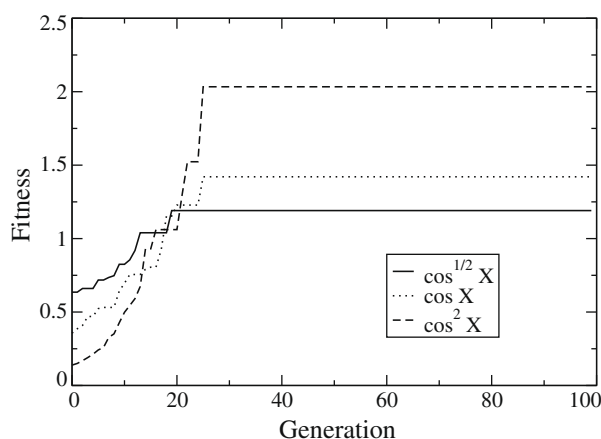

Fig. 4. Fitness functions comparison for the *best_query*, the one for which the greatest precision improvement is achieved.
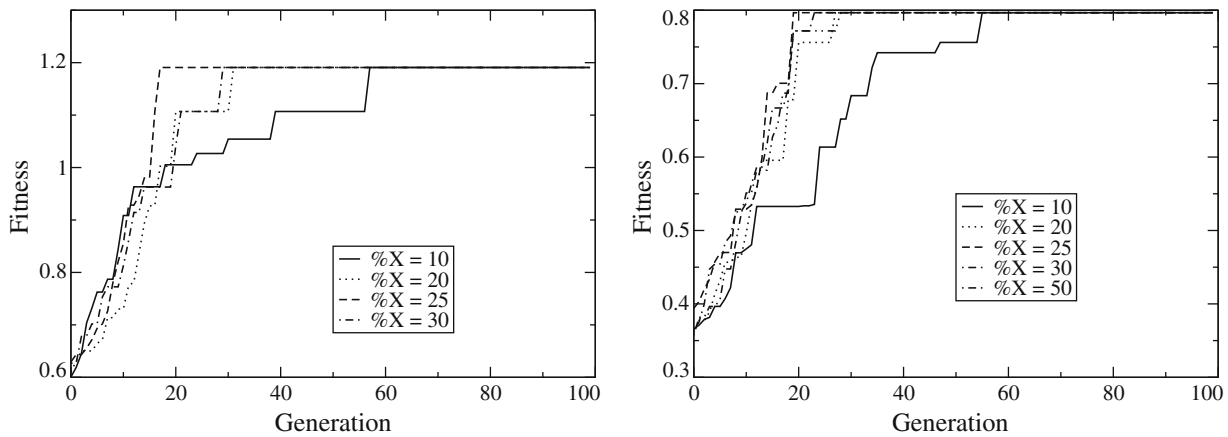
Fig. 5. Studying the best crossover (on the left) and mutation (on the right) rates for the *best_query*. The GA parameters are a population size of 100 individuals, a mutation rate of 1% (in the crossover figure) and a crossover rate of 25% (in the mutation figure).
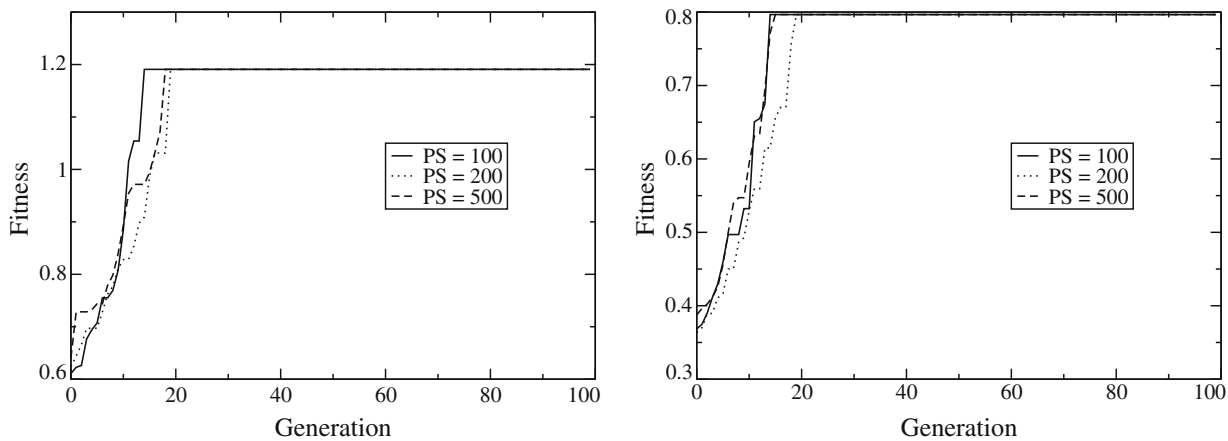

Fig. 6. Studying the population size for the *best_query* (on the left) and the worst one (on the right). The GA parameters are a crossover rate of 25% and a mutation rate of 1%.

Table 4
Stemming Performance. Last column shows the increase achieved with each method with respect to the baseline.

Method	Avg. Prec.	Prec10	Rel. $\varDelta$
No Stemming	.37	.47	−13.9
Stemming (Baseline)	**.43**	**.52**	0
Stem Expansion	.20	.28	−53.5
Stem Expansion (Clauses)	**.43**	**.52**	0
Genetic Expansion	.39	.49	−9.4
Genetic Expansion (Clauses)	**.45**	**.53**	+4.4

stemming (*Porter Stemming*). We can observe in Table 4 (Genetic Expansion (Clauses)) that the combination of clauses and genetic algorithm achieved an improvement of the performance, greater than the one achieved with other stemming methods traditionally used in the stemming process, such as Porter.

5. Related works

In most query expansion literature terms are selected (globally from the entire corpus or locally from the top retrieved documents), weighted with respect to their potential relevance and then passed on to a standard retrieval system, which is considered a black box. Here we are concerned only with this black box and not with the expansion process; for this reason we will not review the query expansion literature (an up to date overview can be found in [5]). Some work on user and pseudo-feedback has tackled the issue of term re-weighting, from early Rochio algorithms to more modern probabilistic approaches of relevance feedback. While these works discuss the ranking function, to our knowledge they all assume Query Term Independence and concentrate on the re-weighting formula. Again, we are not concerned here on the re-weighting of terms (this is left unspecified in our work), and therefore we do not review this literature further (see for example [4]).

A few papers have dealt with the issue of term *correlation* and its effect on retrieval. In [27] the problem of correlation is discussed in depth. They remark that *term correlation* is only an abstract concept and can be understood in a number of ways. They measure term correlation in terms of *term co-occurrence*. Furthermore they propose to represent documents not in the space of terms but in the space of *minterms* which are sets of highly correlated terms. This has the effect of *decorrelating* the terms in the query with respect to hypothetical *concepts* (formally defined as minterms). Instead of computing all term correlations, Silva et al. [24] propose to mine association rules to compute the most significant term correlations and the rotates the term vectors to reflect the extracted correlations; this yields a more selective term de-correlation. Pôssas et al. [20] also propose mining association rules to find term sets of correlated terms. However, the ranking function adjustment proposed is based on the same idea of this paper: collapsing term frequencies within a clause. In fact, if we disregard relative weights, we use the VSM model, and we construct query clauses using association rules in [20], the ranking function here is exactly the same as in [20]. However our work differs from the previously cited papers in that it is not tied to an extraction method or a ranking model, it does not specify the form of the term correlations and furthermore it assumes that term correlations will be *query-dependant*.

6. Conclusions and future work

In this paper we try to show the importance of term dependence issues, how they show up unexpectedly in simple experiments and how they can have a strong adverse effect in performance. Furthermore we propose a method to represent and take into account a simple form of dependence between terms.

On the other hand, we have shown how the clauses can be combined with an evolutionary algorithm to help to reformulate a user query to improve the results of the corresponding search. Our method does not require any user supervision. Specifically, we have obtained the candidate terms to reformulate the query from a *morphological thesaurus*, with provides, after applying stemming, the different forms (plural and grammatical declinations) that a word can adopt. The evolutionary algorithm is in charge of selecting the appropriate combination of terms for the new claused query. To do this, the algorithm uses as fitness function a measure of the proximity between the query terms selected in the considered individual and the top ranked documents retrieved with these terms.

We have investigated different proximity measures as fitness functions without user supervision, such as cosine, square cosine, and square-root cosine.

We have also studied the GA parameters, and see that small values such as a population size of 100 individuals, a crossover rate of 25% and a mutation rate of 1%, are enough to reach convergence. Measures on the whole test set of queries have revealed a clear improvement of the performance, both over the baseline, and over other stemming expansion methods.

A study of the queries resulting after the reformulation has shown that in many cases the GA is able to add terms which improve the system performance, and in some cases in which the query expansion worsen the results, the GA is able to recover the original query.

For the future, we plan to investigate the use of other sources of candidate terms to generate the claused queries applying different query expansion approaches like co-occurrence measures or methods based in Information Theory.

Acknowledgements

References

[1] Lourdes Araujo, Jose R. Pérez-Agüera, Improving query expansion with stemming terms: a new genetic algorithm approach, in: EvoCOP, 2008.
[2] Ricardo A. Baeza-Yates, Berthier A. Ribeiro-Neto, Modern Information Retrieval, ACM Press/Addison-Wesley, 1999.
[3] James P. Callan, W. Bruce Croft, Stephen M. Harding, The inquery retrieval system, in: DEXA, 1992, pp. 78–83.
[4] Claudio Carpineto, Renato de Mori, Giovanni Romano, Brigitte Bigi, An information-theoretic approach to automatic query expansion, ACM Trans. Inform. Syst. 19 (1) (2001) 1–27.
[5] Youjin Chang, Iadh Ounis, Minkoo Kim, Query reformulation using automatically generated query concepts from a document space, Inform. Process. Manage. 42 (2) (2006) 453–468.
[6] Hsinchun Chen, Ganesan Shankaranarayanan, Linlin She, Anand Iyer, A machine learning approach to inductive query by examples: an experiment using relevance feedback, id3, genetic algorithms, and simulated annealing, JASIS 49 (8) (1998) 693–705.
[7] O. Cordón, E. Herrera-Viedma, M. Luque, Improving the learning of boolean queries by means of a multiobjective iqbe evolutionary algorithm, Inform. Process. Manage. 42 (3) (2006) 615–632.
[8] Oscar Cordón, Félix de Moya Anegón, Carmen Zarco, A new evolutionary algorithm combining simulated annealing and genetic programming for relevance feedback in fuzzy information retrieval systems, Soft Comput. 6 (5) (2002) 308–319.
[9] Oscar Cordón, Enrique Herrera-Viedma, Cristina López-Pujalte, María Luque, Carmen Zarco, A review on the application of evolutionary computation to information retrieval, Int. J. Approx. Reason. 34 (2–3) (2003) 241–264.
[10] Er Gelbukh, Mikhail Alex, Sangyong Han, Detecting inflection patterns in natural language by minimization of morphological model, in: Proceedings of 9th Iberoamerican Congress on Progress in Pattern Recognition, Image Analysis and Applications, CIARP 04, Lecture Notes in Computer Science, vol. 3287, Springer-Verlag, 2004, pp. 432–438.
[11] David E. Goldberg, Genetic Algorithms in Search Optimization and Machine Learning, Addison Wesley, 1989.
[12] John J. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, 1975.
[13] Jorng-Tzong Horng, Ching-Chang Yeh, Applying genetic algorithms to query optimization in document retrieval, Inform. Process. Manage. 36 (5) (2000) 737–759.

[14] Cristina Lopez-Pujalte, Vicente P. Guerrero Bote, Félix de Moya Anegón, A test of genetic algorithms in relevance feedback, Inform. Process. Manage. 38 (6) (2002) 793–805.
[15] M. Smith, M. Smith, The use of genetic programming to build boolean queries for text retrieval through relevance feedback, J. Inform. Sci. 23 (6) (1997) 423–431.
[16] José Luis, Fernández-Villacañas Martín, Mark Shackleton, Investigation of the importance of the genotype–phenotype mapping in information retrieval, Future Gener. Comput. Syst. 19 (1) (2003) 55–68.
[17] Zbigniew Michalewicz, Genetic Algorithms + Data Structures = Evolution Programs, second ed., Springer-Verlag, 1994.
[18] Carol Peters, Martin Braschler, European research letter: cross-language system evaluation. The clef campaigns, JASIST 52 (12) (2001) 1067–1072.
[19] Frederick E. Petry, Bill P. Buckles, Thyagarajan Sadasivan, Donald H. Kraft, The use of genetic programming to build queries for information retrieval, in: International Conference on Evolutionary Computation, 1994, pp. 468–473.
[20] Bruno Pôssas, Nivio Ziviani Jr., Wagner Meira, Berthier Ribeiro-Neto, Set-based vector model: an efficient approach for correlation-based ranking, ACM Trans. Inform. Syst. 23 (4) (2005) 397–429.
[21] Alexander M. Robertson, Peter Willet, An upperbound to the performance of ranked-output searching: optimal weighting of query terms using a genetic algorithm, J. Doc. 52 (4) (1996) 405–420.
[22] S.E. Robertson, S. Walker, Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval, in: SIGIR'94: Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, New York, NY, USA, Springer-Verlag, New York, 1994, pp. 232–241.
[23] E. Sanchez, H. Miyano, J. Brachet, Optimization of fuzzy queries with genetic algorithms. application to a data base of patents in biomedical engineering, in: VI IFSA Congress II, 1995, pp. 293–296.
[24] Ilmerio R. Silva, Joao Nunes Souza, Karina S. Santos, Dependence among terms in vector space model, in: IDEAS'04: Proceedings of the International Database Engineering and Applications Symposium (IDEAS'04), IEEE Computer Society, Washington, DC, USA, 2004, pp. 97–102.
[25] Lynda Tamine, Claude Chrisment, Mohand Boughanem, Multiple query evaluation based on an enhanced genetic algorithm, Inform. Process. Manage. 39 (2) (2003) 215–231.
[26] C.J. van Rijsbergen, Information Retrieval, second ed., Butterworths, London, 1979.
[27] S.K. Michael Wong, Wojciech Ziarko, Vijay V. Raghavan, P.C.N. Wong, On modeling of information retrieval concepts in vector space, ACM Trans. Database Syst. 12 (2) (1987) 299–321.
[28] Jing-Jye Yang, Robert R. Korfhage, Query modification using genetic algorithms in vector space models, Int. J. Expert Syst. 7 (2) (1994) 165–191.

Lourdes Araujo received her B.S. degree in Physics and Ph.D. in Computer Science in 1987 and 1994, respectively, both at Universidad Complutense de Madrid (Spain).

After spending three years in Alcatel working on the design of communication networks, she joined the Universidad Complutense in 1990, where she has been teaching Computer Science for 16 years. She is currently Computer Science Professor at the National Distance Learning University (UNED) of Spain and belongs to the Natural Language Processing and Information Retrieval Group of this University.

Her current research interests include natural language processing and information retrieval as well as evolutionary algorithms. She has taken part in several research projects concerning evolutionary algorithms and its application to natural language processing. In particular, she has applied this technique to part-of-speech tagging, parsing and noun phrase detection.

Hugo Zaragoza leads the Natural Language Retrieval group in Barcelona. He is interested in the applications of machine learning (ML) and natural language processing (NLP) for information retrieval (IR), Web search, and information access in general.

From 2001 to 2006, he worked at Microsoft Research (Cambridge, UK) with Stephen Robertson, mostly on probabilistic ranking methods for corporate and web search, but also on document classification, expert finding, relevance feedback, and dialogue generation for games. He also collaborated with Microsoft product groups MSN-Search and SharePoint Portal Server.

Prior to Microsoft Research, he taught computer science and completed a Ph.D. at the LIP6 (U. Paris 6) on the application of dynamic probabilistic models to a wide range of Information Access problems.

José Ramón **Pérez-Agüera** is a researcher at the Metadata Research Center of the School of Information and Library Science in the University of North Carolina. His research interests include metadata, ontologies, taxonomies and thesaurus, as well as different information retrieval aspects, such as probabilistic models, language models, and evaluation techniques.

Prior to this position he has taught computer science and completed two Ph.D.s at the Universidad Complutense de Madrid, one of then in Documentation, and the other one in Computer Science, this one focused on query expansion techniques. He has spent six months in Yahoo! Research Barcelona. Has several publications in international journals and conferences and is member of the advisory board for the distribution list IWETEL and of the working group Thinkepi.

Joaquín Pérez-Iglesias obtained his Computer Science degree by the Rey Juan Carlos University (Madrid) in 2003. As a student he was collaborating with the Departamento de Lenguajes y Sistemas Informáticos at Rey Juan Carlos University. In 2007 he obtained a M.Sc. in Computer Science at the Universidad Complutense Madrid. Simultaneously he has been working with different companies as IBM at Manchester UK and Atos Origin R&D at Madrid.

Currently he is teaching Assistant and Ph.D. candidate at the Departamento de Lenguajes y Sistemas Informaticos at the Universidad Nacional de Educación a Distancia (UNED), in Madrid, Spain, and a member of the research group Natural Language Processing and Information Retrieval.

His primary research interest are related with Query Performance Prediction, Information Extraction, and Natural Language Processing application to Information Retrieval.