

Experiments on Statistical Approaches to Compensate for Limited Linguistic Resources

G.M. Di Nunzio, N. Ferro, and N. Orio

Department of Information Engineering, University of Padua
{dinunzio, nf76,orio}@dei.unipd.it

Abstract. Information Retrieval systems can benefit from advanced linguistic resources when carrying out tasks such as word-stemming or query translation. The main goal of our experiments has been the development of methodologies that minimize the human labor needed for creating linguistic resources for new languages. For this purpose, we have applied statistical techniques to extract information directly from the collections.

1 Introduction

We participated in the monolingual of CLEF 2004 track using a Hidden Markov Model approach for stemmer generation. Our main objective was to study the problem of managing two completely unknown and, for us, non-understandable languages like Finnish and Russian. Building on our experience gained in the campaign of the previous year, we also participated in the bilingual track (German-to-French) to test a new solution for the problem of query expansion/translation from one language to another, especially when linguistic resources are low.

2 Monolingual Track Experiments

The main goal of our monolingual experiments has been the development of methodologies and techniques that do not require, or minimize, the human labor needed when applying information retrieval (IR) techniques to new languages. In this respect, languages such as Finnish and Russian are particularly suitable because they are very different from the languages known by the members of our research group – i.e. Italian, French, and obviously English. French can be considered as a reference language for comparing the system performances.

We focused our attention on the development and test of stemming algorithms, the component of an information retrieval system (IRS) which is most related to the structure of a given language. With the goal of minimizing manual work, we continued in this evaluation campaign the development of a set of stemmers based on Hidden Markov Models (HMMs). According to our approach, HMMs do not require any previous knowledge about the morphology of the language to be stemmed and can be trained simply using a set of words automatically extracted from the test collection.

2.1 Hidden Markov Models to Generate Words

Hidden Markov Models (HMMs) are finite-state automata where transitions between states are ruled by probability functions [5]. At each transition, the new state emits a symbol with a given probability. HMMs are called *hidden* because states cannot be directly observed; only the symbols they emit can be observed. The parameters that completely define an HMM are, for each state: the probability of being the initial or the final state, the probabilities of transition to any other state, and the probability of emitting a given symbol.

HMMs are particularly suitable to model processes that are unknown but can be observed through a sequence of symbols. For instance, the sequence of letters that forms a word in a given language can be considered as a sequence of symbols emitted by an HMM. The HMM starts in an initial state and performs a sequence of transitions between states emitting a new letter at each transition, until it stops in a final state. In general, several state sequences, or *paths*, can correspond to a single word. It is possible to compute the probability of each path, and hence to compute the most probable path corresponding to a word. This problem is addressed as *decoding* and solved using the Viterbi algorithm.

In order to apply HMMs to the stemming problem, a sequence of letters that forms a word can be considered the result of a concatenation of two subsequences: a prefix and a suffix. A way to model this process is through an HMM where states are divided in two disjoint sets: states in the *stem-set* generate the first part of the word and states in the *suffix-set* generate the last part, if the word has a suffix. For many Indo-European languages, there are some assumptions that can be made on the model: 1) Initial states belong only to the stem-set – a word always starts with a stem; 2) Transitions from states of the suffix-set to states of the stem-set always have a null probability – a word can be only a concatenation of a stem and a suffix; 3) Final states belong to both sets – a stem can have a number of different derivations, but it may also have no suffix.

Once a complete HMM is available for a given language, stemming can be carried out straightforwardly considering a word as a sequence of symbols emitted by the HMM. As a first step, the most probable path that corresponds to the observed word is computed using decoding. The analysis of this path highlights the transition from a state of the stem-set to a state of the suffix-set. We call this transition the *split-point*. If there is no split-point then the word has no suffix, otherwise the sequence of letters observed before the split-point is taken as the stem and the one observed after it is taken as the suffix.

2.2 Training the HMM

The proposed topology defines the number of states, their labels indicating the sets to which they belong, the initial and final states, and the allowable transitions. Yet all the probability functions that constitute the HMM parameters need to be computed. The computation of these parameters is normally achieved through *training*, which is based on the Expectation-Maximization (EM) [5] algorithm. It is important to stress that our goal is to develop fully automatic stemmers that do not require previous manual work. This means that we con-

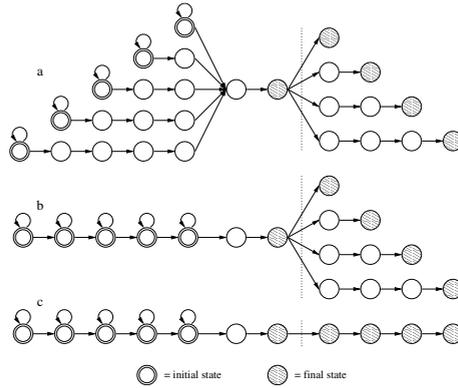


Fig. 1. Three topologies of the HMM that have been tested

sider that neither a formalization of morphological rules nor a training set of manually stemmed words is available.

We propose to perform an unsupervised training of the HMM using only a sample of the words of the considered language. The training set can be built at random by documents that are available at indexing time. It can be noted that an unsupervised training does not guarantee that the split-point of the most probable path has a direct relationship with the stem and the suffix of a given word. In order to create such a relationship, we add some more knowledge about the general rules for the creation of word inflections. Thus, we make the reasonable assumption that, for each language, the number of different suffixes is limited compared to the number of different stems. Suffixes are a set of letter sequences that can be modelled by chains of states of the HMM. This assumption suggests a particular topology for the states in the suffix-set, which can be made by a number of state chains with different lengths, where: transitions from the stem-set are allowed only to the first state of each chain; the transition from one state to the next has probability one; each chain terminates with a final state. The maximum length of state chains gives the maximum length of a possible suffix. Analogously, also the stem-set topology can be modelled by a number of state chains, with the difference that a state can have non-zero self-transition probability. The minimum length of a chain gives the minimum length of a stem. Some examples of topologies for the suffix-set are depicted in Figure 1, where the maximum length of a suffix is set to four letters, and the minimum length of a stem is set to three letters.

After the redefinition of the suffix-set topology, the HMM can be trained using the EM algorithm on a set of words. Given the previous assumption, it is likely that a sequence of letters that corresponds to a suffix will be frequently present in the training set. For this reason, the EM algorithm will assign a higher probability of emitting the letters of frequent suffixes to the states in the suffix-set. For example, considering the suffix-set chains, the state in the one-state chain will emit the last letter of each word with the highest probability, the

states in the two-states chain will respectively emit the most frequent couple of ending letters of each word, and so on. Once the model has been trained, for each word the corresponding path that terminates with the most frequent sequence of letters is expected to have a high probability of being selected as the most probable path, giving a correct stemming of the word.

The STON Algorithm. We developed an algorithm, named STON, to test the methodology and the variations in retrieval effectiveness depending on some of its parameters. STON receives as input a sequence of letters corresponding to a word and gives as output the position of the split-point. As explained in the previous sections, STON needs to be trained off-line using a subset of the words of the collection. Stemming can then be carried out on-line for any new word, also for words that are not present in the training set. The algorithm can be divided in three main steps:

1. *Training/off-line*: given a set of words $w \in W_L$, taken from a collection of documents written in a language L , and an HMM with parameters λ which define the number of states and the set of allowable transitions, STON computes through the Expectation-Maximization algorithm: $\lambda_L^* = \arg \max_{\lambda} \prod_{w \in W_L} Pr(w | \lambda)$.
2. *Decoding/on-line*: given a word w_L , written in language L , and a trained model λ_L , STON computes the most probable path q across the states corresponding to w_L by Viterbi decoding: $q^* = \arg \max_q Pr(q | w_L, \lambda_L)$.
3. *Stemming/on-line*: once the most probable path q^* for word w_L is computed using model λ_L , the split-point can be computed by a simple inspection of the path, i.e. when the state sequence enters the suffix-set.

3 Bilingual Track Experiments

The main goal of our bilingual experiments was to test the effectiveness of an IRS when advanced tools for query or document translations are not available. This situation applies each time an existing IRS has to be extended to a set of new languages, when it is not possible to acquire reliable translators for each language couple – or for each language from/to a single language (e.g. English). It may even be the case that advanced translators are completely unavailable from/to languages that are spoken by a reduced amount of the world population – or that are spoken in countries where economic and technological growth is still slow. We argue that an IRS should have reasonable performance even when linguistic resources are minimal.

In our approach we only considered simple word-by-word translations, such as the ones provided by most of the free translation services on the Web, as tools for the bilingual experiments in the evaluation campaign. As is well known, word-by-word translations have a number of drawbacks, mainly due to the absence of a context for word disambiguation and because of the need to deal with synonyms and antonyms. Clearly these drawbacks can have a negative impact

on the performance of a bilingual IRS. In order to obtain a base translator, we used the Web translation service offered by the Google Web search engine, using the translation of single words in order to not take advantage of the possible use of linguistic cues by the Google translator. It should be noted that this choice imposed an additional constraint, there was no control on the size of the vocabulary.

The methodology that we propose to partially overcome the problems arising from a simple word-to-word translator is based on the use of two collections, the one in the language of the topic and the one in the language of the relevant documents. The source collection is used to expand the query terms that are to be translated. The methodology is presented in the following sections.

3.1 Almost Comparable Corpora

There has been extensive research on the combination of documents collections written in different languages for bilingual information retrieval. Usually the assumption is that two collections of documents, written in two different languages, allow the coupling of documents. The two collections are normally referred to as *parallel corpora* [4] when they can be exactly aligned. This may be the case of transcriptions of legal documents for bilingual countries such as Canada, or when one collection is the translation (made by human experts) of the other one. The two collections are normally referred as *comparable corpora* [6] when documents are independently written in the two languages, but it is possible to couple a subset of the documents, which have the same topic. The coupling can be done using external information such as metadata. An example of this situation is a bi- or multilingual newspaper, where news articles that are potentially of interest for the whole population are independently written in all of the languages, while more local news is reported in articles written in only one language; the GIRT collection, which is available for this evaluation campaign, can be considered an example of comparable corpora.

Unfortunately, parallel and/or comparable corpora are not available for all language couples. Yet there are a number of collections written in different languages that, though not really comparable, have documents that share similar subjects. As an example, we can consider the collections of news articles used for the CLEF campaign. Newspapers in Europe are independently written, yet there are a number of events, for instance of political, social, or economic nature, that are of interest for most European citizens. These events are likely to be the subject of a set of news articles in all of the different newspapers. Clearly, the number of news articles that each newspaper dedicates to a given subject may dramatically vary, depending on political choices, on the locality of the event, and so on. Also the time span in which a given subject is treated can be different. In any case, it is likely that important events give rise to *threads* of news articles, and that threads in different languages can be coupled. We refer to this situation as *almost comparable corpora*.

3.2 Automatic Thread Identification

We propose to use news threads in almost comparable corpora to improve the performance of a bilingual IRS. In particular, if we assume that a topic has been the subject of threads in both the source and the target languages, the automatic identification of a thread in the former can help to retrieve relevant documents in the latter.

The first step in our methodology hence regards the automatic identification of news threads in the source language. We propose to apply *hierarchical clustering* [3] to the documents retrieved by querying a monolingual IRS using the topic of interest. To reduce computational load, clustering is applied only to the first K retrieved documents, that is the ones which are potentially more relevant to the topic. The distance measure used to highlight clusters is based on the classic $tf \times idf$ weighting scheme and computed using the cosine of the angle between retrieved documents. The *inverse document frequency* idf used to calculate the distance between documents is not the same as the original idf computed in the first retrieval. The new inverse document frequency is computed on the first K documents chosen for the query expansion. With this approach we try to discover those words less frequent in the first K documents that may give an added value to the ones of the original query.

The clustering step gives a partition of the set of retrieved documents. We made the assumption that clusters are strictly correlated to news threads. The choice of the most relevant threads for the topic can be based on different strategies. For instance, threads with highest average rank can be chosen, as well as threads that contain the documents with the top 5 or top 10 documents. In our experiments, we chose to select only one thread, the one that contained the document with the highest rank, and to stop the clustering step when the number of clusters is 10 or the present distance between documents belonging to different clusters is less than 0.9. This choice allows us to get a good agglomeration of documents without forcing documents or clusters that are distant, in the sense of the cosine angle, for each other to merge together. On the other hand, it is possible that the cluster that contains the document with the highest rank consists of only that document.

3.3 Topic Translation

Once the potentially most relevant thread in the source language has been highlighted, it can be used to retrieve potentially relevant documents in the target language. We are still investigating the possibility of directly coupling threads in the source and the target languages. For the purpose of the present evaluation campaign, our interest was more focused on the use of thread identification to partially overcome the drawbacks of simple word-by-word translators.

We thus applied query expansion techniques to the topic in the source language, by enriching the bag of words of the topic with a set of words taken from the highlighted thread. In particular, we chose to add words that were more discriminating of the thread, that is the ones with high average tf inside the thread and high idf inside the set of retrieved documents. It is likely that this

additional set of words gives a more complete description of the topic of interest, and also that it contains synonyms of the words in the topic that may allow for a more effective translation.

This expanded set of words can be translated word-by-word, possibly applying stemming, to obtain a set of tokens to be used by an IRS in the target language. The resulting ranked list of documents can be taken as the final result of a run on a given topic. Moreover, the same principle of thread identification can be applied to the target language, in order to rerank the documents according to the threads to which they are potentially relevant. Threads are then highlighted in both source and target language, a coupling between threads might help to refine the results of the run. These steps have not been investigated in detail yet, and are left as future work.

3.4 The Thread-Based Algorithm

According to the methodology outlined above, we developed an algorithm that has been implemented in our IRS. The main step can be described as follows.

1. The topic in the source language is used to query an IRS in the same language. A ranked list of potentially relevant documents is obtained. At this step stemming can be carried out on the source language to possibly improve the performance of the IRS.
2. The first K documents, where K is an integer that has been tested within the range $50 \leq K \leq 100$, are used as the initial set of singletons for hierarchical clustering. The distance between pairs of documents is computed as the cosine of the angle between documents.
3. The merging of clusters stops when the number of created clusters is 10 or when the current distance between documents is 0.9. The cluster that contains the top ranked document is taken as the news thread that is most likely to be relevant to the topic in the source language.
4. The H words that are good candidates to improve query translation are used to expand the initial query in the source language, obtaining an extended set of tokens. A maximum of 10 words for the title of the query and 100 words for the description of the query have been used to expand queries during tests. The algorithm for choosing these words is described in the following section.
5. All the tokens in the extended set are translated singly, using an on-line translation Web service, obtaining an extended set of translated tokens. These tokens are added to the word-by-word translation of the topic in the source language. We used the Google translation Web service.
6. All of the tokens in the target language are used to query an IRS in the same target language. Also at this step, stemming can be carried out in order to improve the system performances.

Finding Good Candidate Words of a Cluster. Once the news *thread* has been found in the first K retrieved documents, how do we choose the words that

are going to expand the original query? We propose the following approach: we use a $(tf \times idf)$ -like weighting scheme to weight terms of a cluster. In particular, we consider the documents of the cluster as boolean vectors, thus we do not make an explicit use of term frequency. For each word w we count how many documents contain w and name it *cluster frequency CF*, and use it in a $(tf \times idf)$ -like formula.

The following steps describe the algorithm used to choose the terms of the cluster. The algorithm is repeated both for the words of the title of a document and the words of the body of a document. In this way, we can add specific title terms to the title of the query, and specific body terms to the description or narrative of a query.

1. Get the cluster that contains the first ranked document in the original retrieval; let N be the number of documents of the cluster, and n be the number of documents in which the word w appears.
2. Let the *Cluster Frequency CF* of w be equal to n , and let *relative Cluster Frequency rCF* be n/N .
3. Calculate the weight of each word as the product $(1 + \log(CF)) \times (-\log(rCF))$.
4. At this point, cluster words according to their weights, that is to say create sets of words with the same weights and order these sets by decreasing order.
5. Add the first set of terms to the query; continue to add sets of words until the number of words that has been added exceeds a predefined threshold.

4 Experiments

We ran a set of experiments in both the monolingual and bilingual tracks using a prototype system that has been developed in our research group. The experimental information retrieval system, called IRON (Information Retrieval ON) was used for the first time in the CLEF 2002 evaluation campaign; it was completely re-engineered for CLEF 2003, and has been further expanded for CLEF 2004 to support different character sets and to include the proposed approach for news thread identification.

IRON is a java multi-threaded program, which provides IR functionalities and enables concurrent indexing and searching of document collections for both monolingual and bilingual tracks. It provides a modular environment suitable for testing the performance of different IR components, and allowing us to easily plug-in the components under examination, such as lexical analyzers (lexers) or stemmers, at runtime. IRON consists of the following components: a Lexer, an IR engine, a Monolingual Track Manager, a Bilingual Track Manager, and a Logger. IRON is partnered with two other tools: WebIRON, a Java Servlet based Web interface, and IRON-SAT, a Matlab program that interacts with IRON in order to carry the statistical analysis of the experimental results.

A description of the experimental results is presented in the following sections.

4.1 Monolingual Experiments

The aim of the experiments for the monolingual track was to compare the retrieval effectiveness of the language independent stemmer, illustrated in Section 2, with that of an algorithm based on a-priori linguistic knowledge – we have chosen the widely used Porter’s stemmers. The hypothesis was that the proposed probabilistic approach generates stemmers that perform as effectively as Porter’s stemmers. In order to evaluate stemming algorithms, we have compared the performances of different IR systems by only changing the stemming algorithms for different runs, all other things being equal. Our aim was to test the following hypotheses: H' – stemming does not hurt and can enhance the effectiveness of retrieval, H'' – the proposed statistical stemmers perform as effectively as those of Porter. Experiments were conducted for the following languages: Finnish, French, and Russian. For each track the following stemming algorithms were tested:

- **No Stem**: no stemming algorithm was applied;
- **Porter**: the stemming algorithms freely available at the Snowball Web site edited by Martin Porter for different languages have been used. Russian is an exception, because the java implementation for the Russian stemmer seems not to properly process Unicode strings and so we were unable to produce runs with this stemmer;
- **STON**: the stemming algorithm based on HMMs has been used.

Tables 1 and 1 report the general figures for the 2004 monolingual topics. Both these tables show that stemming improves all the performances figures for all the considered languages. These figures thus give a positive answer to both hypotheses H' and H'' because stemming improves the performance of an IRS. Moreover, the experimental evidence confirms the hypothesis that it is possible to generate stemmers using probabilistic models without or with very little knowledge about the language. However the degree to which the observed differences are significant has to be measured using statistical testing methods. We used the Wilcoxon signed ranks test [2], which is a non parametric statistical test for paired samples. The runs have been compared query-by-query using the same figures reported in Table 1 with a significance level $\alpha = 5\%$.

Tables 2 and 3 allow us to answer question H' for both Porter and STON. For Finnish, the Porter stemmer exhibits an impact on the performances for all the considered measures; STON shows significant differences with respect to the case of no stemming in terms of number of relevant retrieved documents, but not for the other measures. For French, both stemmers show significant differences with respect to the case of no stemming in terms of number of relevant retrieved documents and average precision. Finally for Russian, STON shows significant differences with respect to the case of no stemming in terms of number of relevant retrieved documents, but not for the other measures. Thus, in general, the hypothesis that stemming influences the performances of an IRS cannot be rejected. The impact of the stemming depends on both the language and the measure considered . Table 4 allows us to answer to hypothesis H'' for the

Table 1. General figures for 2004 monolingual topics

(a) Relevant retrieved document number (recall)

Algorithm	Relevant Retrieved (Recall %)		
	Finnish	French	Russian
No Stem	258 (62.46)	763 (83.38)	82 (66.66)
STON	305 (73.84)	809 (88.41)	94 (76.42)
Porter	346 (83.77)	832 (90.92)	–
Total Relevant Docs	413	915	123

(b) Precision

Algorithm	Average Precision (%)			Exact R-Precision (%)		
	Finnish	French	Russian	Finnish	French	Russian
No Stem	39.62	38.64	28.40	36.12	38.64	28.45
STON	40.70	41.53	34.06	36.75	39.55	30.56
Porter	46.31	42.53	–	43.71	38.71	–

Table 2. Comparison of monolingual No Stem and Porter runs for different measures

Measure		Finnish	French	Russian
Rel. Retr.	No Stem > Porter	0	1	–
	No Stem = Porter	34	33	–
	No Stem < Porter	11	15	–
	Signed Rank Test (p -value)	0.10%	0.08%	–
Avg. Prec.	No Stem > Porter	12	14	–
	No Stem = Porter	10	5	–
	No Stem < Porter	23	30	–
	Signed Rank Test (p -value)	3.06%	0.53%	–
Exact R-Prec.	No Stem > Porter	7	8	–
	No Stem = Porter	22	28	–
	No Stem < Porter	16	13	–
	Signed Rank Test (p -value)	4.97%	49.79%	–

STON algorithm. The results show that in general the hypothesis that STON is as effective as Porter’s algorithm cannot be rejected. However, for Finnish and French there are significant differences between STON and Porter’s stemmers in terms of number of relevant retrieved documents, where Porter’s algorithm performed better than STON.

4.2 Bilingual Experiments

As soon as we received the results for the CLEF bilingual track, we discovered an anomalous behavior of the system. In particular, some of the translated queries were completely empty. We found the error in the following point: when a word is translated from German to French, an apostrophe may appear. Since IRON

Table 3. Comparison of monolingual No Stem and STON runs for different measures

Measure		Finnish	French	Russian
Rel. Retr.	No Stem > STON	1	3	0
	No Stem = STON	35	34	25
	No Stem < STON	9	12	9
	Signed Rank Test (p -value)	1.37%	0.54%	0.39%
Avg. Prec.	No Stem > STON	17	14	11
	No Stem = STON	10	5	7
	No Stem < STON	18	30	16
	Signed Rank Test (p -value)	65.83%	2.15%	6.79%
Exact R-Prec.	No Stem > STON	8	8	2
	No Stem = STON	25	25	27
	No Stem < STON	12	16	5
	Signed Rank Test (p -value)	56.28%	20.86%	57.81%

Table 4. Comparison of monolingual STON and Porter runs for different measures

		Finnish	French	Russian
Rel. Retr.	STON > Porter	2	1	–
	STON = Porter	36	40	–
	STON < Porter	7	8	–
	Signed Rank Test (p -value)	1.95%	2.73%	–
Avg. Prec.	STON > Porter	15	20	–
	STON = Porter	9	5	–
	STON < Porter	21	24	–
	Signed Rank Test (p -value)	11.44%	19.92%	–
Exact R-Prec.	STON > Porter	7	11	–
	STON = Porter	24	31	–
	STON < Porter	14	7	–
	Signed Rank Test (p -value)	10.59%	51.35%	–

reads a query with the `StreamTokenizer` Java class, we discovered that this class contains a method (namely `quoteChar()`) that is used when a quote character is encountered. If a string quote character (in our case the apostrophe) is encountered, then a *quotation* string is recognized, consisting of all characters after the string quote character, up to the next occurrence of that same string quote character, or a line terminator, or end of file. This causes IRON to discard parts of (or even whole) queries and consequently to perform badly (around 13% of average precision). When the error was fixed the general performances of the system improved significantly (up to 23% of average precision).

Table 5 presents the correct results of the bilingual track. As can be noted, this time the differences between stem and no-stem is very subtle. Also the difference between using the first 50 or the first 100 documents is small, although the runs using the first 50 documents obtain a better recall than those using the

Table 5. General figures for bilingual German \rightarrow French

Algorithm	Recall %	Avg. Prec. (%)	Exact R-Prec. (%)
No Stem, first 50 docs	61.85	22.15	21.38
No Stem, first 100 docs	56.50	22.25	21.57
STON, first 50 docs	61.20	22.50	22.36
STON, first 100 docs	53.11	22.41	21.92
Porter, first 50 docs	63.38	23.17	23.00
Porter, first 100 docs	53.22	22.83	22.73
Total Relevant Docs	915	–	–

first 100 documents. We performed the same statistical analysis as described in the previous section, but it does not give any significant difference between stem and no stem, and between using 50 or 100 documents for clustering. There is an exception to this general trend: Porter, first 50 docs, shows a better average precision ($p = 3.57\%$) than No Stem, first 50 docs.

5 Conclusions and Future Work

The idea of minimizing human labor and computing resources was the main point of the many experiments that were carried out this year by the IMS research group. The automatic stemmer generation using Hidden Markov Models was confirmed as a valid alternative to language dependent stemmers such as Porter's. Statistical tests on the monolingual track gave evidence against rejecting the hypothesis that, in general, STON is as effective as Porter's algorithm and that stemming does not hurt and can even enhance retrieval performances. Automatic identification of news threads together with hierarchical clustering was used for query expansion/translation for the bilingual track, results are encouraging. Further experiments and refinements may make performances obtained with this approach comparable to state-of-the-art systems.

References

1. Di Nunzio, G. M., Ferro, N., Melucci, M. and Orio, N.: The University of Padova at CLEF 2003: Experiments to Evaluate Probabilistic Models for Automatic Stemmer Generation and Query Word Translation. In Peters, C., ed.: Working Notes for the CLEF 2003 Workshop, pp. 211–223, 2003.
2. Gibbons, J. D.: Nonparametric Statistical Inference. Marcel Dekker, Inc., New York, USA, 2nd edition, 1985.
3. Johnson, S. C.: Hierarchical Clustering Schemes. *Psychometrika*, 32, pp. 241–254, 1967.
4. Nie, J.-Y., Simard, M., Isabelle, P. and Durand, R.: Cross-language Information Retrieval Based on Parallel Texts and Automatic Mining of Parallel Texts from the Web. Proc. of the 22nd ACM SIGIR Conference, pp. 74–81, Berkeley, CA, 1999.

5. Rabiner, L. and Juang, B. H.: Fundamentals of speech recognition. Prentice Hall, Englewood Cliffs, NJ, pp. 321–389, 1993.
6. Sheridan, P. and Ballerini, J. P.: Experiments in Multilingual Information Retrieval Using the SPIDER System. Proc. of the 19th ACM SIGIR Conference, pp. 58–65, Zurich, Switzerland, 1996.