Query Generation Techniques for Patent Prior-Art Search in Multiple Languages

Dong Zhou, Jianxun Liu, and Sanrong Zhang

Key Laboratory of Knowledge Processing and Networked Manufacturing &
School of Computer Science and Engineering,
Hunan University of Science and Technology,
Xiangtan, Hunan 411201, China
{dongzhou1979,wuyunzsr}@hotmail.com, ljx529@gmail.com

Abstract. Patent prior-art search is an necessary step to ensure that no previous similar disclosures were made before granting an patent. The task is to identify all relevant information which may invalidate the originality of a claim of a patent application. Using the whole patent or extracting high indicative terms to form a query reduces the search burden on the user. To date, There are no large-scale experiments conducted specifically for evaluating query generation techniques used in patent prior-art search in multiple languages. In the following paper, we firstly introduced seven methods for generating patent queries for ranking. Then a large-scale experimental evaluation was carried out on the CLEF-IP 2009 multilingual dataset in English, French and German. A detail comparison of the different methods in terms of performance and efficiency has been performed in addition to the use of full-length documents as queries in the patent search. The results show that some methods, work well in information retrieval in general, fail to achieve the same effectiveness in the patent search. Different methods demonstrated distinct performance w.r.t query and document languages.

Keywords: Patent Prior-Art Search, Multilingual Information Access, Query Generation.

1 Introduction

Patent information retrieval is an active sub-domain of information retrieval that aims to support patent experts to retrieve patents that satisfy their information needs and search criteria [1]. A common scenario in patent information retrieval is prior-art search, which is performed by patent experts to ensure that no previous similar disclosures were made before granting a patent. They will normally need to use some sort of information retrieval systems and tools to automate this process. To improve the usefulness of such systems and tools, researchers from around the world gathered in CLEF[1] and NTCIR[2] for automating this specific

[1] http://www.clef-initiative.eu/
[2] http://research.nii.ac.jp/ntcir/index-en.html

G. Zhou et al. (Eds.): NLPCC 2013, CCIS 400, pp. 310–321, 2013.

task. Many state-of-art patent search systems have been developed. A general trend is to use the full patent document as an input, after preprocessing, search over previous filed patents with the aim of retrieving relevant documents, which may invalidate or at least describe prior art work in a patent application [2–4]. This is of high commercial value to many companies and organizations.

The aim and challenges of patent prior-art search are different from those of standard ad-hoc information retrieval and/or web search. One challenge would be the vocabulary mismatch between existing filed patents and the query patent. This is often caused by the patent writing style. Long patent queries comprising of several hundreds of terms fail to represent a focused information need required for high precision retrieval. On the other hand, the primary focus of the patent prior-art search is to retrieve all relevant documents at early ranks. Carefully balance of precision and recall would be necessary.

Recent work in both CLEF and NTCIR favor to use either full patent documents as queries or key terms extracted from a patent application to produce a more focus information need. For example, all participants in CLEF 2010 [5] and CLEF 2011 [3] adopted the same way. TF-IDF scheme, language model-based weighting scheme, text summarization or phrases extraction techniques were frequently used [6–9]. Multilingual aspect has been addressed in recent CLEF campaigns, specifically for English, French and German [5, 3]. The tasks organized in those workshops did not restrict the language used for retrieving the documents, but participants were encouraged to use the multilingual characteristic of the collection. This is because the claims in granted patent documents may be provided in all three languages. Researchers attempted different search tasks on the provided data. However, there is lack of direct comparison between the performance in different languages.

In this paper, we firstly introduce seven methods for generating query representations. These include a method for removal of unit frequency terms, the TF method, the TFIDF method, the BM25 method, the language model-based approach, the relevance feedback-based method, and the method based on IPC classification. Large-scale experimental evaluation was then carried out on the CLEF-IP 2009 dataset investigating retrieval effectiveness across different languages used. This includes English, French and German. A detail comparison of the different methods in terms of performance and efficiency has been performed in addition to the use of full-length documents as queries in the patent search. Note that the scenario investigated in this paper can not be treated as cross-lingual or multilingual patent search [10, 11], as it only deals with monolingual retrieval in three different languages.

The remainder of this paper is organized as follows. In §2, we summarise related work from the fields of patent information retrieval. In §3, we describe seven methods for generating patent queries. §4 documents the experiments we used to evaluate the methods in three different languages. §5 presents our results. §6 concludes the paper and proposes future work.

2 Related Work

Our work relates to patent information retrieval in general. It is an active research field [3, 5]. To date, a significant amount of development is driven by the Intellectual Property task within the CLEF initiative and NTCIR workshops. The systems developed early days at these evaluation campaigns replicated the work performed by patent examiners, who consider high term frequency in the document to be strong indicator of a good query term [12, 13]. A recent line of work advocates the use of full patent application or automatically extracted terms as the query to reduce the burden on patent examiners. Xue and Croft [4] firstly conducted a series of experiments in order to examine the effect of different patent fields in a patent document. Their work on the USPTO corpus concluded that the best Mean Average Precision is achieved using the query generated from description section of the query patent with raw term frequencies. However, the relevance judgements in their system were not annotated by real patent experts but rather automatically extracted from the citation fields.

Terms extracted from description field have been proved to produce highest retrieval performance by many other research teams. For example, Magdy et al. [14] showed that the second best performing run of CLEF-IP 2010 uses a list of citations extracted from the patent numbers within the description field of some patent queries. Mahdabi et al. also confirmed in a series of experiments that under a language model framework, terms extracted from the description field shown to be effective [2, 15]. They also showed that automatically disambiguated query terms could be informative by extraction of noun phrases from the global analysis of the patent collection. However, the use of phrases is of some controversial. Becks et al. [7] demonstrated that with a different patent corpus (CLEF 2011 rather than CLEF 2010) phrase queries reported negative results.

Another important feature of the patent retrieval w.r.t to the ordinary ad-hoc search and web search is that the pseudo-relevance feedback technique performs poorly in this particular context [16, 2]. This may be due to the reasons that the precision at top ranks is usually low, the information focus is not clear and added terms are noisy. Ganguly et al. [16] tackled this problem from a different angle. They decomposed a patent application into constituent text segments and computed the language similarities by calculating the probability of generating each segment from the top ranked documents. However, their method could be viewed as a very slight modification to use the query with a full set of terms in a document.

Due to the increase and distribution of inventions across the world, the necessity of research and commercial tools that support patent search in different languages has increased. Despite the popularity of query generation techniques described above, to our knowledge, there is no comparison between using these techniques in multiple languages.

3 Query Generation Techniques

The reference model presented in our paper uses the full length document as input after stopwords removal, stemming and number removal. Unlike many previous experiments conducted by different CLEF participants [7, 17], we did not use patent-specific stopwords, phrases or sophisticated summarization methods. This model is denoted as $FULL$ subsequently.

Clearly the $FULL$ method fails to represent a focused information need required for high precision retrieval. We now introduce seven different query generation techniques that use many state-of-art methods to reduce the full-length query. The following three models extract key terms according to the usual TF-IDF, BM25 and simple term frequency schemes. Defined as follows:

$$P(t|q_{TFIDF}) = \frac{n(t,d)}{\max_{t'} n(t',d)} \cdot \log \frac{N}{df_t}$$

$$P(t|q_{TF}) = \frac{n(t,d)}{\max_{t'} n(t',d)}$$

$$P(t|q_{BM25}) = \sum_t w_t \frac{(k_1+1)n(t,d)}{K+n(t,d)} \frac{(K_3+1)n(t,d)}{k_3+n(t,d)}$$

where $w_t = \log \frac{N-df_t+0.5}{df_t+0.5}$, $K = k_1 \cdot ((1-b) + b \cdot \frac{|d|}{avg|d|})$, and $n(t,d)$ is the term frequency, df is document frequency. These three query generation techniques were denoted as $TFIDF$, TF and $BM25$ respectively.

The next technique adopts the unigram model proposed by Mahdabi et al. [2], which could be regarded as a strong baseline because it produces comparable results to the second best runs in CLEF 2010. It is defined by estimating the importance of each term according to a weighted log-likelihood based approach as expressed below:

$$P(t|q_{LM}) = Z_t P(t|\Theta_q) \log \frac{P(t|\Theta_q)}{P(t|\Theta_C)}$$

where Θ_q and Θ_C are language model estimation for a term in the query patent and in a test collection, respectively. $P(t|\Theta_q)$ is defined as:

$$P(t|\Theta_q) = (1-\lambda) \cdot P_{ML}(t|d) + \lambda \cdot P_{ML}(t|C)$$

with $P_{ML}(t|d) = \frac{n(t,d)}{\sum_{t'} n(t',d)}$. $Z_t = 1/\sum_t P(t|q)$ is the normalization factor and defined as the Kullback-Leibler divergence between Θ_q and Θ_C. This model is named LM in the reminder of the paper.

The next technique moves one step further by considering International Patent Classifications (IPC[3]) in the patent documents as in [2]. IPC provides for a hierarchical system of language independent symbols for the classification of patents and utility models according to the different areas of technology to which

[3] `http://www.wipo.int/classifications/ipc/en`

they pertain. Thus the IPC classes resemble tags assigned to documents. We build a relevance model Θ_{IPC} by employing documents that have at least one tag in common with the query topic. The result model is defined as:

$$P(t|q_{IPC}) = (1 - \lambda) \cdot P(t|\Theta_{IPC}) + \lambda \cdot P(t|q_{LM})$$

where $P(t|\Theta_{IPC})$ is calculated by using:

$$P(t|\Theta_{IPC}) = \sum_{d \in IPC} P(t|d) \cdot P(d|\Theta_{IPC})$$

and

$$P(D|\Theta_{IPC}) = Z_d \sum_{t} P(t|\Theta_d) \log \frac{P(t|\Theta_{IPC})}{P(t|\Theta_C)}$$

where $Z_d = 1/\sum_{D \in IPC} P(D|\Theta_{IPC})$ is a document specific normalization factor. This query generation technique is denoted as $LMIPC$ subsequently.

The sixth technique we used is simply full-length query by removing unit frequency terms (i.e. terms which occur only once in the patent query), denoted as UFT henceforth.

The last technique, denoted as QR, is a method using Pseduo Relevance Feedback for reducing patent queries [16]. The technique decomposes a patent application document into s constituent text segments (sentences) and computes the language modeling similarities by calculating the probability of generating each segment from r top ranked document. Finally, the method selects τ fraction of sentences to retain in the query. The language similarity equation is shown below:

$$logP(q_s|d) = \sum_{t \in q_s} n(t) log(1 + \frac{\lambda P(t|d)}{(1 - \lambda)} P(t))$$

The final score of each segment is the sum of all r documents. This completes our description of the query generation techniques used in our experiments. In the next section we will detail a large-scale experiment involving read-world patent data in three different languages.

4 Evaluation

In the following section, we describe a series of experiments designed to answer the following questions:

1. How effective are the state-of-art query generation techniques when used in patent prior-art search
2. Does the query produced by the seven query generation techniques perform better than using the full-length document as the query?
3. How effective are the various techniques in different languages?
4. What is the optimal number of keywords for selective-based query generation techniques (i.e. $BM25$, TF, $TFIDF$, LM and $LMIPC$)?

4.1 Experimental Data

To make a fair comparison, the text corpus used in our evaluation was built using components of the CLEF-IP 2009 test collections. This collection contains patents, physically stored as a collection of XML files encoding patent document. A patent document maybe an application document, a search report, or a granted patent document. The data is extracted from the MAREC [4] data corpus and contains a number of approximately 1,958,955 million patent documents, referring to approximately 1,022,388 million patents. We also used the CLEF-IP 2009 query set, which contains 311 English topics, 164 German topics and 25 French topics. Each topic is a patent application composed of several fields (e.g. *Title, Abstract, Description, Claims* etc.). We used the *Description* field for building the query as it previously showed the best performance. We used relevance judgements produced by the CLEF workshops. Note that we did not use the citation information of the patent applications in our experiments. Prior to indexing and retrieval, a suffix stemmer [18] and a stopword list[5] were applied to all documents and queries for English texts. We did not apply any linguistic processes on German and French texts.

4.2 Evaluation Metrics

We used the following evaluation metrics in this experiment:

– The precision of the top 10, top 50 and top 100 documents (P@10, P@50 and P@100)
– Normalized Discounted Cumulative Gain (NDCG) [19]
– The recall of the top 10, top 50 and top 100 documents (R@10, R@50 and R@100)
– Mean average precision (MAP).

Unless otherwise stated, the results given indicate *average performance* across all test topics. Statistically-significant differences in performance were determined using a paired t-test at a confidence level of 95%.

4.3 Retrieval Systems

All information retrieval functions in our experiments were handled by the Terrier open source platform[6] [20]. As described in §3, the results are obtained by seven models: $TFIDF$, TF, $BM25$, LM, $LMIPC$, UFT and QR. For comparison purposed we also used $FULL$ as our reference model.

[4] It is a collection of over 19 million patent documents, available from information retrieval facility (`http://www.ir-facility.org/`)

[5] `ftp://ftp.cs.cornell.edu/pub/smart/`

[6] `http://terrier.org/`

4.4 Parameter Settings

The parameters in the baseline systems are set according to the tuning procedures in their original papers if detailed. k_1, b, k_3 used in the $BM25$ model were set to 1.2, 0.75 and 7 respectively. The number of pseduo-relevant documents (i.e. r) used in QR were set to 20, s and τ were set to 20 and 90% respectively.

5 Results

5.1 Precision-Oriented Performance

In our first evaluation, we compare the precision-oriented performance across all methods described. Statistical significant results are obtained when using $FULL$ as the baseline. The number of key terms selected for $BM25$, TF, $TFIDF$,

Table 1. Precision-oriented performance across three languages, statistically significant results are marked with $*$

French					
	MAP	**NDCG**	**P@10**	**P@50**	**P@100**
FULL	0.0435	0.1031	0.048	0.0128	0.008
BM25	0.0695*	0.1426*	0.064*	0.0152*	0.0084*
UFT	0.0389	0.0879	0.04	0.0112	0.006
TF	0.0465	0.098	0.048	0.0112	0.006
TFIDF	0.0559*	0.1195*	0.052	0.0128	0.008
LMIPC	0.044	0.111	0.044	0.012	0.008
LM	0.0547	0.1173	0.048	0.0136	0.0084
QR	0.0432	0.1013	0.052	0.0128	0.008
German					
	MAP	**NDCG**	**P@10**	**P@50**	**P@100**
FULL	0.0732	0.164	0.0524	0.0163	0.0103
BM25	0.1045*	0.2155*	0.0707*	0.023*	0.013*
UFT	0.0716	0.1565	0.0543	0.016	0.0098
TF	0.0745	0.1648	0.0524	0.0162	0.0099
TFIDF	0.0907*	0.1921*	0.0622*	0.0191*	0.0122*
LMIPC	0.0753	0.1687	0.0549	0.0183*	0.0113
LM	0.0864*	0.1875*	0.0598*	0.019*	0.012*
QR	0.0715	0.162	0.0506	0.0162	0.0098
English					
	MAP	**NDCG**	**P@10**	**P@50**	**P@100**
FULL	0.0825	0.2105	0.0717	0.0289	0.0178
BM25	0.0794	0.2109	0.0688	0.0282	0.0175
LMIPC	0.0893*	0.2312*	0.0746	0.0301	0.0196*
LM	0.0881	0.2306*	0.0768*	0.03	0.0188
QR	0.0808	0.2071	0.0736	0.0284	0.0175
TF	0.0826	0.2226*	0.0736	0.0292	0.0183
TFIDF	0.0903*	0.2319*	0.0781*	0.0306*	0.0196*
UFT	0.0822	0.206	0.0669	0.0257	0.015

LM and $LMIPC$ will be discussed in section 5.3. From Table 1, we note that $FULL$, UFT and QR methods consistently deliver lowest performance, only slightly better than $BM25$ in English. These three methods are quite similar as they keep most of terms in the query. Next, we observed that the TF method only increases the mean average precision by a modest amount. This shows significantly different behaviour when deployed in the USPTO corpus as in [4]. The $TFIDF$ method works well in English and the $BM25$ method works well in French and German. Amazingly, the $BM25$ method demonstrates opposite performance in English test collection, achieved lowest performance. The reason may be the language complexity of French and German comparing to English. Given the performance obtained by $BM25$ and $TFIDF$, we must be cautious in drawing conclusions from the experiments w.r.t to the best term weighting methods in a multilingual setting.

Mahdabi et al. [2] showed that using IPC information could increase the effectiveness of the retrieval system. It is confirmed from the results that using the IPC information could help improving the performance (in terms of $LMIPC$ works better than LM) in English. However, the improvements are only modest and not statistically significant. The same results could not be replicated in French and German. This shows the accuracy of IPC information attached to the French and German documents is low.

Overall, we notice that $TFIDF$ works the best in the English test collection, showing statistically significant improvements across all evaluation metrics. $BM25$ works the best in the French and German test collections, with statistically significant results observed in all metrics.

5.2 Recall-Oriented Performance

We now measure the performance of the methods using various recall-based metrics in three languages. Bare in mind that patent prior-art search is a recall oriented task where the primary focus is to retrieve all relevant documents at early ranks in contrast to ad hoc and web search. Same trend can be observed from Table 2 as in the precision-oriented evaluation. $TFIDF$ works better in English, while $BM25$ works better in French and German. The improvements over $FULL$

Table 2. Recall-oriented performance across three languages, statistically significant results are marked with $*$

	French			**German**			**English**		
	R@10	**R@50**	**R@100**	**R@10**	**R@50**	**R@100**	**R@10**	**R@50**	**R@100**
FULL	0.0789	0.1122	0.1424	0.1116	0.1729	0.2125	0.1151	0.2194	0.2622
BM25	0.1189*	0.1436*	0.1436	0.1554*	0.2346*	0.2621*	0.1093	0.211	0.2557
UFT	0.0556	0.08	0.088	0.1136	0.1636	0.1988	0.1109	0.2021	0.2344
TF	0.0778	0.0889	0.0933	0.1074	0.1672	0.2047	0.1139	0.22*	0.276*
TFIDF	0.0844*	0.1089	0.1391	0.132*	0.1987*	0.246*	0.1273*	0.2309*	0.2783*
LMIPC	0.0833	0.1091	0.1391	0.1135	0.1864	0.2274	0.1247*	0.2366*	0.2961*
LM	0.0778	0.1167	0.1469	0.1268	0.196*	0.2409*	0.1234	0.2295	0.2828
QR	0.0856	0.1122	0.1424	0.1077	0.1714	0.2039	0.117	0.2164	0.2573

were quite stable across all evaluation metrics. $LMIPC$ also frequently delivers statistically significant results. This performance of the IPC-based method has an intuitive explanation. Relevant documents are being 'found' by using the IPC information. The low performance of QR confirms that using PRF is not a wise choice for the patent search. In summary, the methods achieves good precision-oriented performance usually perform well in the recall-oriented evaluation.

5.3 Selection of the Number of Key Terms

Recall that in five of the methods, $BM25$, TF, $TFIDF$, LM and $LMIPC$, the top terms with higher weights must be picked and used to build the query. In this section we study the optimal number across three different languages. Results in Figure 1-3 show that 65-95 terms in English, 20-65 terms in French and 55-90 terms are sufficient to capture the most information in a long query.

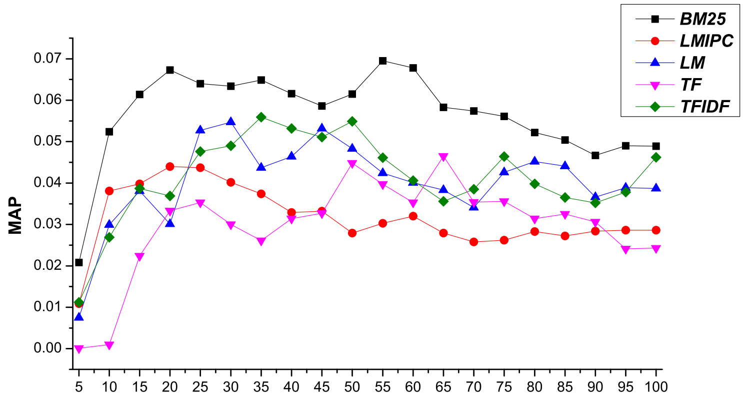

Fig. 1. Selection of the number of key terms in French

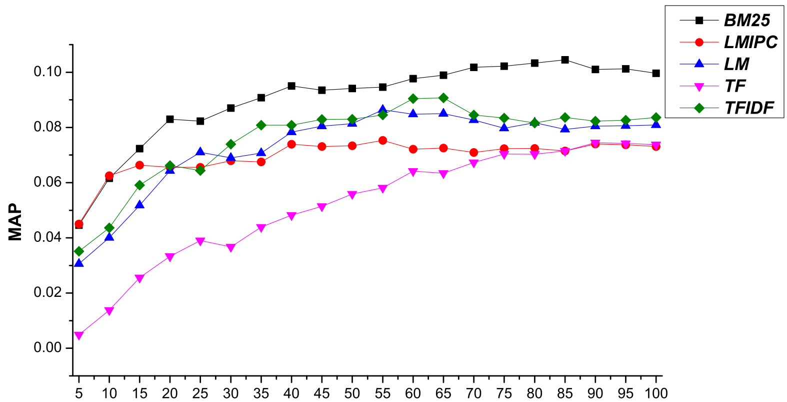

Fig. 2. Selection of the number of key terms in German

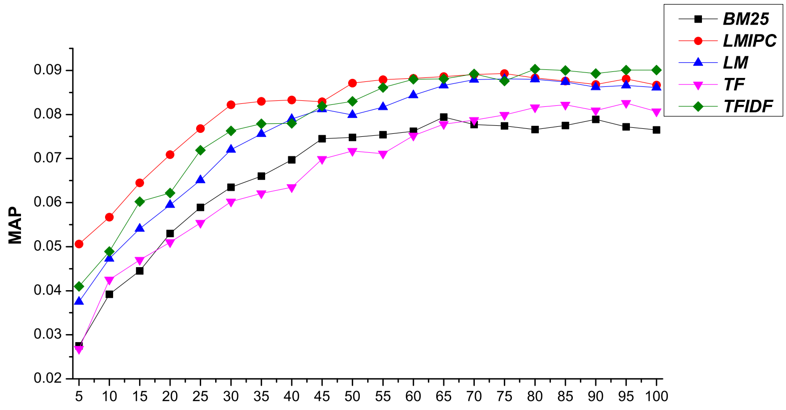

Fig. 3. Selection of the number of key terms in English

6 Conclusion and Further Work

In this paper, we introduce seven methods for generating queries from a full-length patent query. Large-scale experimental evaluation has been carried out on the CLEF-IP 2009 multilingual dataset in English, French and German. A detail comparison of the different methods in terms of performance and efficiency has been performed in addition to the use of full-length documents as queries in the patent search. The experimental results show that the TFIDF method achieved the highest performance in English, the BM25 method works the best in French and German. The methods achieve good precision-oriented performance usually perform well in the recall-oriented evaluation. The paper also found that in general less than 100 selected key terms can obtain good results for selective-based methods.

In future work, we plan to explore additional sources of intellectual property documents beyond CLEF-IP (NTCIR and USPTO) to investigate the differences. We also plan to explore more term weighting methods commonly used in the information retrieval field.

Acknowledgments. This research is supported by the National Natural Science Foundation of China under grant No. 61300129 and No. 61272063, Excellent Youth Foundation of Hunan Scientific Committee No. 11JJ1011.

References

1. Tait, J. (ed.): Proceedings of the 1st ACM workshop on Patent Information Retrieval, PaIR 2008, Napa Valley, California, USA, October 30. ACM (2008)
2. Mahdabi, P., Andersson, L., Keikha, M., Crestani, F.: Automatic refinement of patent queries using concept importance predictors. In: Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2012, pp. 505–514. ACM, New York (2012)

3. Piroi, F., Lupu, M., Hanbury, A., Zenz, V.: Clef-ip 2011: Retrieval in the intellectual property domain. In: CLEF (Notebook Papers/Labs/Workshop) (2011)
4. Xue, X., Croft, W.B.: Transforming patents into prior-art queries. In: Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2009, pp. 808–809. ACM, New York (2009)
5. Piroi, F.: Clef-ip 2010: Retrieval experiments in the intellectual property domain. In: CLEF (Notebook Papers/LABs/Workshops) (2010)
6. Mahdabi, P., Andersson, L., Hanbury, A., Crestani, F.: Report on the clef-ip 2011 experiments: Exploring patent summarization. In: CLEF (Notebook Papers/Labs/Workshop) (2011)
7. Becks, D., Eibl, M., Jürgens, J., Kürsten, J., Wilhelm, T., Womser-Hacker, C.: Does patent ir profit from linguistics or maximum query length? In: CLEF (Notebook Papers/Labs/Workshop) (2011)
8. Magdy, W., Jones, G.J.F.: Applying the kiss principle for the clef- ip 2010 prior art candidate patent search task. In: CLEF (Notebook Papers/LABs/Workshops) (2010)
9. Lopez, P., Romary, L.: Experiments with citation mining and key-term extraction for prior art search. In: CLEF (Notebook Papers/LABs/Workshops) (2010)
10. Zhou, D., Truran, M., Brailsford, T., Wade, V., Ashman, H.: Translation techniques in cross-language information retrieval. ACM Computing Surveys 45(1), 1:1–1:44 (2012)
11. Zhou, D., Truran, M., Brailsford, T., Ashman, H.: A hybrid technique for english-chinese cross language information retrieval. ACM Transactions on Asian Language Information Processing 7(2), 5:1–5:35 (2008)
12. Iwayama, M., Fujii, A., Kando, N., Takano, A.: Overview of patent retrieval task at ntcir-3. In: Proceedings of the ACL-2003 Workshop on Patent Corpus Processing, PATENT 2003, vol. 20, pp. 24–32. Association for Computational Linguistics, Stroudsburg (2003)
13. Itoh, H., Mano, H., Ogawa, Y.: Term distillation in patent retrieval. In: Proceedings of the ACL-2003 Workshop on Patent Corpus Processing, PATENT 2003, vol. 20, pp. 41–45. Association for Computational Linguistics, Stroudsburg (2003)
14. Magdy, W., Lopez, P., Jones, G.J.F.: Simple vs. Sophisticated approaches for patent prior-art search. In: Clough, P., Foley, C., Gurrin, C., Jones, G.J.F., Kraaij, W., Lee, H., Mudoch, V. (eds.) ECIR 2011. LNCS, vol. 6611, pp. 725–728. Springer, Heidelberg (2011)
15. Mahdabi, P., Keikha, M., Gerani, S., Landoni, M., Crestani, F.: Building queries for prior-art search. In: Hanbury, A., Rauber, A., de Vries, A.P. (eds.) IRFC 2011. LNCS, vol. 6653, pp. 3–15. Springer, Heidelberg (2011)
16. Ganguly, D., Leveling, J., Magdy, W., Jones, G.J.: Patent query reduction using pseudo relevance feedback. In: Proceedings of the 20th ACM International Conference on Information and Knowledge Management, CIKM 2011, pp. 1953–1956. ACM, New York (2011)
17. Mahdabi, P., Andersson, L., Hanbury, A., Crestani, F.: Report on the clef-ip 2011 experiments: Exploring patent summarization. In: CLEF (Notebook Papers/Labs/Workshop) (2011)

18. Porter, M.F.: Readings in information retrieval, pp. 313–316. Morgan Kaufmann Publishers Inc., San Francisco (1997)
19. Järvelin, K., Kekäläinen, J.: Ir evaluation methods for retrieving highly relevant documents. In: Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2000, pp. 41–48. ACM, New York (2000)
20. Ounis, I., Amati, G., Plachouras, V., He, B., Macdonald, C., Lioma, C.: Terrier: A High Performance and Scalable Information Retrieval Platform. In: Proceedings of ACM SIGIR 2006 Workshop on Open Source Information Retrieval, OSIR 2006 (August 2006)