Multilingual Patent Text Retrieval Evaluation: CLEF–IP

Florina Piroi and Allan Hanbury

Abstract The CLEF–IP evaluation lab ran between 2009 and 2013 with a two-fold expressed purpose: (a) to encourage research in the area of patent retrieval with a focus on cross language retrieval, and (b) to provide a large and clean data set of patent related data, in the three main European languages, for experimentation. In its first year, CLEF–IP organized one task only, a text retrieval task that modelled the "Search for Prior Art" done by experts at patent offices. In the following years the types of CLEF–IP tasks broadened to include patent text classification, patent image retrieval and classification, and (formal) structure recognition. With each task, the test collection was extended to accommodate for the additional tasks. In this chapter we overview the evaluation tasks dealing with the textual content of the patents. The Intellectual Property (IP) domain is one where specific expertise is critical, implementing Information Retrieval (IR) approaches to support some of its tasks cannot be done without the use of this domain know-how. Even when such know-how is at hand, retrieval results, in general, do not come close to the expectations of patent experts.

1 Introduction

In a nutshell, patents can be seen as contracts between inventors and governments by which the former can exclude other parties from manufacturing and exploiting an invention without permission from the patent owner. This corresponds to a pessimistic view of the patent system based on a "blocking effect," raising a sequence of issues in the modern world, like, for example, invention fragmentation or failures in securing patent licensing (Galasso and Schankerman 2013). On the more optimistic note, the patent system is viewed as fundamental to the diffusion

F. Piroi (✉) · A. Hanbury
Institute of Information Systems Engineering, TU Wien, Vienna, Austria
e-mail: florina.piroi@tuwien.ac.at; allan.hanbury@tuwien.ac.at

N. Ferro, C. Peters (eds.), *Information Retrieval Evaluation in a Changing World*, The Information Retrieval Series 41,
https://doi.org/10.1007/978-3-030-22948-1_15

of ideas and a key incentive to advancing the technological knowledge of society, which some countries underline in their patent law (Kumagai 2005).

In this chapter we look, first, at how the patent system evolved to its current state (Sect. 2.1), then we look at search and retrieval on patent data (Sect. 2.2). In Sect. 2.3 we outline the main phases of the patent's life-cycle and recount the administrative character of the patent systems. We continue with describing the CLEF–IP test collection and describe the text retrieval tasks that were organized in this evaluation lab in Sects. 3 and 4. We finish with a description of the submissions and a submission scores summary.

2 A Background on Patents

As a main governmental instrument to increase research and development (Galasso and Schankerman 2013), patents are not only an output of R&D activities but also an indicator of the technological competitiveness at national, regional or sectoral levels (Frietsch et al. 2010). In this section we give an abridged account of the origins of the modern patent systems. We, then, explain the need for Information Retrieval research in the patent domain, giving an account on the IR research efforts in the IP domain. The section continues with a description of patent data characteristics.

2.1 The Patent System: A Very Brief History

Inventions, as the root of new technologies and developments, provide consistent input to civilization advancements. Until the emergence of the Greek civilization, discoveries and inventive activities were extremely low paced (Skolnik 1977). The first recorded grant of a monopoly refers to the time of the Sybarites (approximately 750 B.C. (Pfaller 2013b) and (Anthon 1841)) when 1 year exclusiveness on exceptional food recipes were awarded.

The emergence of the Greek civilization accelerated the pace of discovery, but the idea of invention was established only by the end of the thirteenth century, at the beginning of the Renaissance (Skolnik 1977). Historians agree that one of the first exclusive rights of use we know about was awarded to a Florentine architect, Fillipo Brunelleschi, for a special type of barge that was capable of transporting heavy loads (marble) along the Arno River (Skolnik 1977; Pfaller 2013a). In the exposure of motives to grant this patent it was shown that the inventor was refusing to reveal his invention for fear that there was not enough protection against others who would replicate and use it. The period of exclusiveness awarded to Brunelleschi was of 3 years.

A few decades later, the first patent system was developed in fifteenth century Venice and was explicitly utilised to promote innovation (May 2010). In March 1474, the Senate of Venice issued a decree which made patents a subject of a generalized law instead of individual petitions and monopoly grants (May 2010; Skolnik 1977).

Other patent laws (in the sense we give it today) were the "Statute of Monopolies," released in 1623, in England, and in 1787, in France, which granted longer periods of exclusive use for inventions (Rich 1993; Skolnik 1977). In America, where the first patent was granted as early as 1641, the first US Patent Act was passed in April 1790, and conferred inventors exclusive rights for 14 years for disclosing their inventions (Skolnik 1977; Mossoff 2007). Later, in 1861, this time period was extended to 17 years. Other European countries also extended and modernized their patent and monopolies laws during the nineteenth century, and during the twentieth century, the use of the patent system became worldwide ubiquitous (Hall 2017).

As national patent systems evolved, the differences in patent laws between nations were considerable. The 1883 Paris Convention for the Protection of Industrial Property successfully established a unified system for multinational filings, enabling worldwide priority to be obtained for an invention originating in any one country part of the treaty (Skolnik 1977; Hall 2017). A 110 years later, in 1995, the adoption of the Trade-Related Aspects of Intellectual Property Rights Agreement (TRIPS), ensures that the patent granting process is approximately the same everywhere in the world (Hall 2017). The World Intellectual Property Organization (WIPO) joined the administrative offices of the Paris Convention of 1883 and the Berne Convention of 1886, which established rules about protection of literary and artistic works. WIPO's[1] first international IP filing service was launched with the adoption of the Madrid Agreement in 1891. In 1978 the Patent Cooperation Treaty (PCT) came into existence, which allows inventors in any of the treaty's signatory countries to file patent applications and seek protection of the invention in countries other than the country of origin (PCT 1970).

Patent rights, when granted, are usually restricted within the border of the patent office jurisdiction. In most of the cases, this is a country, with the exception of the European Patent Office (EPO) and the African Regional Intellectual Property Organization (ARIPO) where the countries of patent jurisdiction are under the control of the applicant (Hall 2017).

[1] Until 1970, what we currently know as WIPO was called *Bureaux Internationaux Réunis pour la Protection de la Propriété Intellectuelle* (French for *The United International Bureaux for the Protection of Intellectual Property*).

2.2 Search and Retrieval on Patent Data

The patent and monopolies offices, in their very early stages, were doing little more than registering, filing, and classifying the inventions.[2] The basic principles of patent examination were laid down by the adoption of the US Patent Act in 1836, principles which were soon adopted by other countries (1902 in France, 1877 in Germany) (Skolnik 1977). Already in the 1970s the information retrieval problem was an issue: with over 600,000 worldwide applications per year, a large number at that time, only partial retrieval solutions were available, most of them based on the classification systems (McDonnell 1969). While studying local clustering in full-text searches using local feedback, experiments were done on a small database of US patents (Attar and Fraenkel 1977). Attar and Fraenkel (1977) did an experiment that was a "technology survey"-like search on a set of 76 US patents. Two decades later a "prior art search" was performed on 13,747 US patents where the topics of the search were patents and citations were used to generate relevance assessments (Osborn et al. 1997).

In the last decades, research in IR methods for the IP domain has intensified. Workshops, conferences and evaluation tracks were organized in an effort to bring IR and IP communities together (see Iwayama et al. 2003; Kando and Leong 2000; Tait et al. 2010; Hanbury et al. 2010). The National Institute of Informatics (NII), Japan, initiated a series of workshops and evaluations using patent data as part of the NTCIR project (the NII Test Collections for IR Systems, currently renamed as the NII Testbeds and Community for Information access Research), focusing on Japanese and Chinese patents, and their translations into English.

In 2009, two further evaluation activities using patent data were launched: TREC–CHEM and CLEF–IP. TREC–CHEM ran from 2009 to 2011 and was organized as a chemical IR track in TREC (Text Retrieval Conference) addressing challenges in Chemical and Patent Information Retrieval (Lupu et al. 2009). The document collection used by TREC–CHEM was limited to chemical patent documents and chemical journal articles.

The purpose of the CLEF–IP track was to encourage and facilitate research in the area of multilingual patent retrieval by providing a large, clean data set for experimentation. The data set contains patents in three main European languages, patents published by the European Patent Office (EPO), as well as queries and associated relevance judgements.

[2]The first US classification system consisted of 16 classes in 1830.

2.3 Characteristics of Patent Data

We give a brief account of the main phases in the patenting process, establishing at the same time the basic patent notions and the different patent related aspects that are used throughout this chapter.

Pre-application Phase A person having developed an invention, will first write down a document describing the invention's background, a detailed description of it, and a set of claims that specify the extent of the protection sought. The level of detail of each of the document parts may vary depending on the patent office. The claims part of the application document is a legal text, therefore it is common to get the help of a patent attorney to draft it. This leads to the patent document having a mixture of writing styles, with the description of the invention being written in a narrative style, while the claims are written in a legal style (also called "attornish" or "patentese").

Before registering this document with a patent office, the inventor usually does what is called a "technology survey" of the existing technology in the area of his or her invention, the results of the search possibly triggering a change in the invention's specifications.

Examination Phase Upon registering the document with a patent office it becomes known as the "Patent Application Document" and receives an alphanumerical code that uniquely identifies it among other patent applications.

When a patent application is filed at a patent office, the application is given to patent professionals for examination. Each patent office follows different laws when deciding which claims to grant, but there is a set of worldwide common criteria that have to be fulfilled by any application before a patent can be granted (EPO 2018):

- **novelty:** the invention should not be previously known;
- **inventive step:** the invention should not be obvious for experts in the technological area of the invention;
- **realizable:** the invention can be manufactured by experts in the area.

The novelty check for an invention is done by performing a thorough search on the data collections available to the patent expert examining the patent application. The novelty search is the most time consuming and expensive part of the application examination. According to personal communications with various patent experts, the examination for novelty can take up to several weeks and even months, searches being repeated sometimes on different areas of the available databases, or with different sets and combinations of keywords. The result of a novelty search (also known as a "Prior Art Search") is a list of relevant documents stored into a "Search Report"; the relevant documents are called patent citations (note the different meaning of the word "citation" compared to academic publications). The citations that an examiner found to be relevant to an application can be of three main types, which (in their order of relevance) are:

- citations that describe prior work but which do not destroy the novelty of the application (lower relevance);
- citations that, in combination with other citations, destroy the novelty of an application;
- citations which, taken alone, make a patent application not novel (high relevance).

Granting and Opposition Phases When the search report is created, a series of official communications between the applicant and the patent office take place. As an output of these communications claims are usually modified in order not to infringe existing patents. Quite often, patent applications are withdrawn.

When the patent office takes the decision to grant a patent, a “Granted Patent Document” is published. From this point on, for a certain amount of time (9 months at the EPO) oppositions to a granted patent may be filed to the patent office.

In this chapter we refer to the documents generated during the patenting phases as “patent documents.” One patent will, administratively, consist of several patent documents, like the Patent Application Document, the Search Report or the International Search Report, the Granted Patent Document.

2.3.1 Types of Patent Search

Depending on the type of information need and on the starting parameters of the search, the process used in finding relevant patents can differ from case to case and from one practitioner to another (Lupu and Hanbury 2013). A detailed description of the types of patent search can be found, for example, in Adams (2011), Alberts et al. (2011), Hunt et al. (2007) and should be differentiated from the IR task of searching in test collections that contain patents.

The search types that are typically performed in the three patent life-cycle phases above are:

- Pre-Application Search (technology survey) which is a search done by the inventor before filing for a patent application. The goal of the search is to identify existing knowledge (printed or not, including patents) which pertains to the invention.
- Novelty Search which aims to establish the novelty or the lack of novelty of an invention. This search can be performed both for filed patent applications or granted patents, as well as for inventions that were not yet filed.
- Patentability or Validity Search which is a search to identify prior art (that is previously published documents) that are relevant to the inventiveness of a patent application. Such searches may include novelty searches and are often carried out during the examination of a patent application.

2.3.2 Patent Data Is Administrative Data

During the patenting process a large number of documents are usually created, both by the patent office and by the applicant or her attorney. Communications to/from the patent office, application document amendments, registration of fee payments, and designating the states where the patent is valid are all examples of information that belong to the patent itself.

The general understanding of the patent concept is that, through its claims, it restricts other parties from exploiting the invention described in the respective granted patent. However, if we view patents as the complete set of documents generated during the patenting process, we immediately notice that patent data has a substantial administrative side. The administrative data includes, for example, application dates, addresses of the inventors and/or patent assignees, priority references, legal status, and so on. Of interest for the CLEF–IP tasks presented here are the patent classification system and the patent families.

Patent Clustering by Families In the current global economy, often enough after filing an initial patent application, inventors will pursue legal protection for their invention in additional countries of interest for them. Following the general patenting process, they will file subsequent applications at each patent office in the countries of interest referring to the original filing as the “priority claim”. Even though these applications may somewhat differ in content, depending on the patent laws in force at the various patent offices, it is obvious that, worldwide, patent content is often replicated. To assist patent practitioners with minimizing the necessary documents they might need to inspect, several methods to group ‘parallel’ patent documents were devised. The group of applications pertaining to the same invention is called a “patent family”.

There is no single definition of what a patent family is. Moreover, each provider of patent data constructs the patent families differently. For example, the EPO uses three types of patent family, while the WIPO additionally defines three further types (WIPO 2013). Nevertheless, as with the patent classification systems, the patent families are widely used when dealing with patent data.

Patent Classification by Technological Areas Patent classification systems are designed to categorize the patent documents by technological areas and sub-areas, using the technical features of the disclosed inventions. Several patent classification systems are in use, systems created both by patent offices and by private companies. The most well known are the International Patent Classification System (IPC),[3] the United States Patent Classification (USPC),[4] the F-term Japanese Classification System (Schneller 2002), or the Derwent Classification System.[5] Since January

[3] International Patent Classification (IPC) www.wipo.int/classifications/ipc/en/.

[4] United States Patent Classification www.uspto.gov/patents/resources/classification/.

[5] Derwent World Patents Index clarivate.com/products/dwpi-reference-center/dwpi-classification-system/.

2013 the EPO and the USPTO (US Patents and Trademarks Office) use a joint classification system, the Cooperative Patent Classification system (CPC).[6]

In the early days of the patent system, patent classification systems were designed as a shelf-location tool for paper files (Adams 2000). Even today, these systems are manually maintained by experts and represent a ubiquitous resource for augmenting the query terms of on-line patent retrieval environments.

3 A Collection of European Patent Documents

One of our aims at the time we embarked on the CLEF–IP endeavour was to create a test collection fit for experimenting with patent data, a collection that faithfully mirrors the features and challenges of the data used in the actual working cycles of a patent professional. For this we use actual patent documents from the EPO and WIPO. These documents contain most of the information that is actively used by patent practitioners in their daily work with patent data.

The bulk of the collection's corpus is made of patent documents stored as XML files. In CLEF–IP, a patent consists of one or more XML files, one for each patent document that was available at the time of the collection creation. Since its first release in 2009, consecutive additions were made to the CLEF–IP test collection, so that it currently contains almost 3.5 million XML files corresponding to almost 1.5 million patents. These patents are an extract from the larger MAREC[7] collection which contains files representing over 19 million patents published at the EPO, USPTO, WIPO and JPO (Japan Patent Office) stored in a common normalized XML format. The main elements of the XML representations are shown in the simplified listing below:

```
 <patent-document>
   <bibliographic-data> ... </bibliographic-data>
   <abstract> ... </abstract>
   <description> ... </description>
   <claims> ... </claims>
</patent-document>
```

The `<abstract>`, `<description>`, and `<claims>` elements store the textual content of the disclosed invention. These fields may occur more than once when, for example, both the English and the German versions of the abstract are stored in a patent document. Most of the patent text retrieval methods make use of the abstract, description and claims fields. The `<bibliographic-data>` element contains the administrative data related to a patent. In this XML element we will find the application and publication dates and references, family identifiers, the patent

[6]Cooperative Patent Classification (CPC) www.cooperativepatentclassification.org/.

[7]The MAtrixware REsearch Collection. http://ifs.tuwien.ac.at/imp/marec.

classification symbols, inventors, assignees, postal addresses of the inventors and/or assignees, the invention title (in three languages), and the patent citations relevant to the invention in this document.

The CLEF–IP collection is limited to the MAREC patents published by the EPO, patents with application date earlier than 2002. The EPO patent documents published later were retained to form a *test and training topic pool* of approximately 500,000 patents, out of which we extracted training sets and topic sets for the CLEF–IP tasks (Graf and Azzopardi 2008).

In the corpus of European patent documents with application date prior to 2002, a high percentage of the patent documents refer to applications internationally filed under the Patent Cooperation Treaty (PCT 1970), also known as “EuroPCTs”. For these filings, the EPO does not republish the whole patent application, but only bibliographic entries that link to the original application published by the WIPO. Using text-based methods to retrieve such documents is problematic, and therefore, for these patent documents we added their WIPO equivalent to the CLEF–IP collection. Determining that the EuroPCT patent documents refer to a certain invention disclosed in a document published by WIPO is done by the patent family identifier which for the two documents must be the same. In this way, the collection became both larger and more realistic.

One of the most important features of the CLEF–IP corpus is its multilingualism. Patent applications to the EPO are written in one of the three official EPO languages (German, English, French), with the additional requirement that, once the decision to grant a patent is made, the claims section of the patent document must be submitted in all these three languages. Although the English language is over-represented in the CLEF–IP collection (see Table 1), not least due to the EuroPCT applications written in their large majority in English, the collection entails large amounts of content that is in German and French, making the collection suitable for carrying out multilingual retrieval experiments.

According to the specifics of each organized task, further chunks of data were added to the core CLEF–IP patent collection. One such data addition consisted of image files occurring in patents intended to support the concurrent use of textual and visual retrieval methods into one multimodal information retrieval method.

Table 1 Document distributions in CLEF–IP

3.1 million documents		
		67% English
14% WIPO documents	74% applications	22% German
86% EPO documents	26% granted patents	6% French
		5% Unknown

4 The CLEF-IP Text Retrieval Tasks

There were five CLEF–IP evaluation cycles with a total of 7 tasks (Table 2). Some of the tasks were organised once only (e.g. the "Chemical Structure Recognition" task), others ran for 2 or 3 years in a row.

The "Prior Art Candidates" task (PAC, 2009–2011) required that, for a given patent application document (the "topic patent"), all patent documents relevant to the described invention are retrieved. The "Passage Retrieval (Starting from Claims)" task (PSG, 2012–2013) required that, given a patent application document and a selected subset of its claims, all patents that may invalidate these claims are retrieved, and, in addition, the concrete passages that do so are returned.

The "Patent Classification" task (CLS, 2010–2011) requested that a given patent document was classified according to the IPC classification symbols.

To solve these three tasks—PAC, PSG, and CLS—only text based analysis of the available CLEF–IP test collection files was necessary. Besides these text retrieval and classification tasks, and as part of the CLEF–IP campaign, further tasks that involved analysis of images in patents were organised between 2011 and 2013.

The "Image-based (Prior Art) Retrieval" task (IMG-PAC, 2011) asked the participants to retrieve relevant patents to the invention in a given topic patent, where, in addition to the text content in the XML patent documents, we provided the images that were attached to the patents. For more details on this task see (Piroi et al. 2011).

The "Image Classification" task (IMG-CLS, 2011) required that 1000 topic patent images (figures attached to patents) were classified into one of nine classes: drawing, chemical structure, program listing, gene sequence, flow chart, graph, mathematics, table, and symbol. No text analysis was necessary for this task.

The "Flowchart/Structure Recognition" task (2012–2013) and the "Chemical Structure Recognition" task (2012) didn't necessitate text analysis either, as they required participants to extract content from patent images and store it into a predefined textual format in order to make it search-able by text-based IR methods.

Table 2 gives an overview of the CLEF–IP tasks and number of topics by the year they were organised. The last four tasks that involve image analysis are not the subject of this chapter, for more details we direct the reader to the references

Table 2 CLEF–IP tasks, number of topics in the main topic sets, and year of their organisation

Task/year	2009	2010	2011	2012	2013
Prior art candidates (PAC)	10,000	2000	3973		
Passage retrieval (PSG)				105	149
Patent classification (CLS)		2000	3000		
Image-based retrieval (IMG-PAC) (Piroi et al. 2011)			211		
Image classification (IMG-CLS) (Piroi et al. 2011)			1000		
Flowchart/structure recognition (Piroi et al. 2012, 2013)				100	747
Chemical structure recognition (Piroi et al. 2012)				865	

```
<topic>
    <num>EP-1222860-A2</num>
    <narr>Find all patents in the collection that
       potentially invalidate patent application
       EP-1222860-A2.</narr>
    <file>EP-1222860-A2.xml</file>
</topic>
```

Fig. 1 Excerpt from the file with the list of topics in CLEF–IP PAC tasks

indicated in Table 2. In the following we detail the design of each CLEF–IP text-related task, the data used to extract topics and relevance judgments for the topics.

4.1 Topic Sources

The topics for each of the PAC, PSG, and CLS tasks consisted of an XML file corresponding to patent applications published between 2002 and 2008, selected from the *test and training topic pool*. In 2009 the topics were selected such that at least one highly relevant patent citation per topic was contained in the CLEF–IP collection. A further condition on topic selection, in 2009, was that, for a topic patent, the XML patent document is a Granted Patent Document which, according to the EPO regulations, provides the claims in the three EPO official languages (German, English, French).[8] With this decision we gave the task participants incentives to investigate cross-language retrieval methods already in the first CLEF–IP evaluation cycle.

In 2010 and 2011, to model the IP professional work procedures and rules more realistically, the topic patents are Patent Application Documents. We sampled the topic patents by their document language, by available citations within the CLEF–IP collection, and by their IPC class, such that each IPC class is equally represented in the final topic test set.[9] To further stimulate the research into cross-language patent retrieval methods, whenever possible, we selected topic patents where the language of the patent citation document was different from the language of the patent application document language (e.g. application document language is English, while the document language of a relevant patent citation in the search report is French or German).

The list of topics is stored as an XML file where the topic identifier is the patent number as assigned by the EPO (Fig. 1).

[8]The occurrence of multi-lingual content is a consequence of the Rule 71(3) of the European Patent Convention (EPO 1973) which states that granted patents must contain claims in the three official languages of the EPO.

[9]IPC classification represents the different domains of the patent applications: chemistry, textiles, mechanical engineering, physics, electricity, etc.

```
<tid>PSG-2</tid>
<tfile>EP-1445439-A1.xml</tfile>
<tfam-docs>FI-116479-B1.xml,FI-20030196-A.xml,FI-20030196-D0.xml</tfam-docs>
<tclaims>/patent-document/claims/claim[1] /patent-document/claims/claim[2]
/patent-document/claims/claim[3] /patent-document/claims/claim[4]</tclaims>
```

Fig. 2 Excerpt from the file with the list of topics in CLEF–IP PSG tasks

Examining the EPO patent search reports closer, we immediately observe that, besides the list of patent citations relevant to a patent application, the reports detail which parts of a citation document (lines, columns, figures, etc.) are pertinent to which particular claims of the patent application. Therefore, in 2012 and 2013, we changed the PAC task formulation from ‘find relevant documents’ to ‘find relevant documents and mark in them the passages of interest to a given set of patent application claims’ (PSG). At the same time, although the basis for topic creation remained the same—actual patent application documents from the topic pool—the topics are now (sub)sets of claims in the patent application document, instead of the patent application document itself. It also allowed us to extract more than one topic (set of claims) out of one patent application document (Piroi et al. 2012, 2013). Figure 2 is an example of a topic in the CLEF–IP 2013 PSG list of topics file: Although the PSG topics contained only claims, it was allowed to use other parts of the topic’s application patent document for query generation. Moreover, in 2013, each topic contained also the reference to the patent document that constituted the priority claim document of the topic application document. Examiners at patent offices also have access to this kind of information related to new, incoming patent applications.

We note that, for each task and each year, the topic sets did not overlap. Similarly, for each of the three tasks and in each year, distinct sets of training topics were provided to the participants.

We conclude this subsection with a few remarks on the topic document’s language. In 2009, in addition to the main topic set where no restrictions on the document’s language were applied, three additional language specific tasks were created, where the topics in each of the three sets were documents in only one of the three EPO official language. In 2010, where no language specific tasks were organised, we did not impose restrictions on the document language when selecting the topics, which resulted in the obvious fact that the document language distribution in the topic set followed the document language distribution in the collection corpus (see Table 1). A consequence of this ‘natural’ language distribution was that methods using distinct algorithms for the different languages to process, index, and search the documents were not easy to qualitatively assess with respect to their language specific methods. We compensated for this in the following years where each third of the topic set contained documents written in one of the official EPO languages. The same is true for the training sets as well, where each EPO language was represented by a third of the topics.

4.2 Relevance Assessments and Metrics

Any organiser of an IR evaluation campaign faces the challenge of how to best obtain the ground truth for the topic test sets in order to be able to judge the quality of the submitted retrieval results. The big majority of the evaluation efforts (TREC, CLEF) use some form of document pooling from the submitted retrieval experiments, manually assessing the relevance of the documents in the pool by volunteer work (Spark-Jones and Van Rijsbergen 1975; Voorhees and Harman 2005). Recently, efficient pooling strategies have been proposed such that human effort may be reduced (Lipani et al. 2017). Still, obtaining humanly created relevance assessments is time-consuming and, in the case of patent evaluation, volunteers are difficult to find as costly expert knowledge is required (Roda et al. 2010). At the same time, because of strict regulations in logging their work, patent experts at patent offices do provide partial relevance assessments in the form of patent citations in the search reports. These relevance assessments are of high quality and, furthermore, at the EPO, the patent citations have relevance degrees assigned to them (see Sect. 2.3, Examination Phase).

However, using search reports as a source for relevance assessments gives an average of six relevant documents for a patent application document. This low number did not change over the years. In 1996/1997, in their experiment with patent retrieval, Osborn et al. found that their test collection also showed an average of six documents per query (Osborn et al. 1997). Nevertheless, we extracted relevance assessments from patent search reports following the general lines described in Graf and Azzopardi (2008). To increase the number of relevant documents we made use of patent families by creating an extended list of citations which includes the patent citations of the topic patent application document, the patent citations of the topic document's family members and the family members of the patent citation documents. After filtering out the patent citations that are not part of the CLEF–IP corpus, we reached an increase in the number of relevant documents by a factor of 7 (Roda et al. 2010).

As explained above, we used patent families to extract relevance assessments for the PAC topics. Obtaining the relevance assessments for the CLS task was straight forward: the IPC relevant classes were extracted from the classification assigned by the patent offices and present in the administrative part of the documents (the `<bibliographic-data>` XML field).

Extracting the relevance assessments for the PAC and the CLS tasks could be done automatically. The relevance files contain lists of ⟨*topic, relevant document*⟩ identifier pairs, where the identifiers referred to documents in the collection. The situation was more challenging for the PSG task, where we could not make use of patent families any more. In this task both the topics and the relevance assessments contain XPaths to the claims and relevant passages in the XML patent documents. The relevance files contain lists of ⟨*topic, relevant document, relevant passage XPath*⟩ identifier triples where the relevant document identifier refers to patent documents relevant to the topic, and the passage XPath identifies, within the relevant

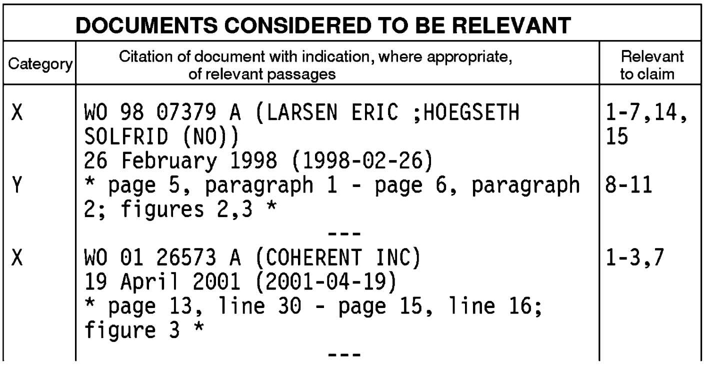

DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim
X	WO 98 07379 A (LARSEN ERIC ;HOEGSETH SOLFRID (NO)) 26 February 1998 (1998-02-26)	1-7,14, 15
Y	* page 5, paragraph 1 - page 6, paragraph 2; figures 2,3 * ---	8-11
X	WO 01 26573 A (COHERENT INC) 19 April 2001 (2001-04-19) * page 13, line 30 - page 15, line 16; figure 3 * ---	1-3,7

Fig. 3 Extract from a search report

document, the passage that is pertinent to the claims in the PSG topic. For the PSG task the relevant passage information was extracted manually by matching the passage indications in the search reports (Fig. 3) with the textual content of the patent documents in our corpus. When matched, we extracted the XPaths of the identified content and saved them to a database. This process was time consuming, the main hurdle being comparing the PDF patent documents to which the search reports refer with the XML content of the document in the CLEF–IP collection. Therefore, the number of topics in the PSG test sets is low compared to the number of topics in the PAC and CLS tasks.

The measures reported for the PAC tasks are Precision and Recall at different cut-offs, MAP, nDCG (Järvelin and Kekäläinen 2002), and PRES (Magdy and Jones 2010a). For the CLS tasks we computed Precision, Recall and F_1 at one classification code and at five classification codes (a patent may be classified into more than one IPC class). Since the PSG relevance assessments were triples, the evaluation for this task could be done on two levels: at the relevant document level and at the relevant passage (XPath) level. The evaluation at the document level measured a system's performance in retrieving whole relevant documents, very similar to the evaluations done in the PAC task, while the evaluation at the passage level targeted measuring the ranking quality of the passages in the relevant patent documents (Piroi et al. 2012). At the document level we maintained the computation of MAP, Recall and PRES measures. At the passage-level we assessed the systems' quality w.r.t. the relevance of the returned passages (XPaths) by computing MAP and Precision scores for the retrieved passages grouped by relevant documents (MAP(D) and Precision(D)) and then averaging over the set of topics. These two document level measures carry similarities with the 'Relevant in Context' metrics of the INEX campaign (Kamps et al. 2008), but looking at sequences of XPaths instead of sequences of characters (Piroi et al. 2012, Section 2.1).

5 Submissions and Results

For all CLEF–IP tasks, a *submission* (or *run*) consisted of a single text file with at most 1000 answers per topic. The most answers were given for the Prior Art Tasks, while the Patent Classification tasks required fewer answers per topic. With few variations, the format of the submissions followed the format used for the TREC submissions, which is a list of tuples containing at least the *topic identifier*, the retrieved *document* (and *passage* for the PSG tasks), the *rank* of the retrieved answer, and the *score* given by the retrieval system to the retrieved answer. Table 3 lists the groups that have submitted experiments to the PAC, PSG, and CLS tasks.

Generally, participants in the CLEF–IP evaluation benchmark have used off-the-shelf retrieval and classification engines (Indri/Lemur or Terrier engines, commonly available k-nearest neighbour algorithm implementations, support vector machines, SVM, or Winnow-like classifiers), choosing to tune these systems on the provided training sets. The better results, however, were obtained by those systems that put more effort into understanding and exploiting the patent specific data, like citations or classification symbols (Lopez and Romary 2009, 2010; Magdy and Jones 2010b; Mahdabi et al. 2011).

Some of the participants did experiments to determine which parts of the (topic) patent documents contribute most to improving retrieval results. These included selecting certain file parts to index, building separate indexes per document XML field, or boosting query terms extracted from certain parts of the topic files (Gobeill et al. 2009; Becks et al. 2010; Gobeill and Ruch 2012; Verberne and D'hondt 2011).

Given that each patent document could contain text in up to three languages, some participants chose to build separate indexes per language (Lopez and Romary 2009; Szarvas et al. 2009), while others generated one mixed–language index or used text fields only in one language discarding information given in the other languages (Correa et al. 2009; Toucedo and Losada 2009). Few participants made use of machine translations to obtain query terms in additional languages and applying them on the previously created collection indexes (Magdy and Jones 2010b). The granularity of the index varied, too, as some participants chose to concatenate all text fields into one index, while others indexed different fields separately. In addition, several specific indexes like phrase or passage indexes, concept indexes and IPC indexes were used (Magdy et al. 2009; Wanagiri and Adriani 2010; Szarvas et al. 2009). A more detailed analysis of the indexing methods and of the retrieval approaches used in the 2009 and 2010 evaluation labs can be found in Piroi and Zenz (2011).

As the task topics were complete documents, with several pages of texts, extracting appropriate queries from the topic document has been investigated by several participating teams (Graf et al. 2009; Becks et al. 2009).

The IPC classification codes were the part of the `<bibliographic-data>` that was exploited the most and was used either as a post-processing filter, as part of the query, or to pre-select smaller sets of patents to search in Gobeill et al. (2009), Szarvas et al. (2009), Eiselt and Oberreuter (2013), Lopez and Romary (2009),

Table 3 Teams that participated in the text-based retrieval and classification CLEF–IP tasks

Team		2009	2010	2011	2012	2013
BiTeM, Service of Medical Informatics, Geneva Univ. Hospitals	CH	PAC	PAC CLS		PSG	
Centrum Wiskunde & Informatica - Interactive Information Access	NL	PAC	PAC CLS			
Chemnitz Univ. of Technology, Dept. of Computer Science	DE			PAC	PSG	
Dublin City Univ., School of Computing	UK	PAC	PAC			
Geneva Univ., Centre Universitaire d'Informatique, SimpleShift	CH	PAC	CLS		PSG	
Gerogetown Univ., Dept. of Computer Science	US					PSG
Glasgow Univ. - IR Group Keith	UK	PAC				
Hewlett-Packard Labs, Russia	RU			PAC		
Humbold Univ., Dept. of German Language and Linguistics	DE	PAC	PAC CLS			
Industrial Property Documentation Dept., JSI Jouve	FR		CLS			
Innovandio S.A.	CL					PSG
Inria	FR	PAC	PAC CLS			
SIEL, International Institute of Information Technology	IN			PSG		
LCI – Institut National des Sciences Appliqu'ees de Lyon	FR		CLS			
Radboud Univ. Nijmegen	NL	PAC	CLS	CLS		
Santiago de Compostela Univ., Dept. Electronica y Computacion	ES	PAC				
Spinque B.V.	NL		PAC CLS	PAC		
Swedish Institute of Computer Science	SE	PAC				
Technical Univ. Darmstadt, Dept. of CS, Ubiquitous Knowledge Processing Lab	DE	PAC				
Technical Univ. Valencia, Natural Language Engineering	ES	PAC				
UNED - E.T.S.I. Informatica, Dpto. Lenguajes y Sistemas Informaticos, Madrid	ES		PAC			
Univ. Indonesia, Information Retrieval Group	ID		PAC			
Univ. "Alexandru Ioan Cuza", Iaşi	RO	PAC	PAC			
Univ. of Hildesheim, Information Science	DE	PAC	PAC	PAC	PSG	
Univ. of Lugano	CH			PAC	PSG	
Univ. of Macedonia, Dept. of Applied Informatics, Thessaloniki	GR				PSG	PSG
Univ. of Neuchatel, Computer Science	CH	PAC				
Univ. of Tampere - Info Studies & Interactive Media	FI	PAC				
Univ. of Wolverhampton, School of Technology	UK				PSG	
Vienna Univ. of Technology, IFS	AT			PAC	PSG	PSG
WISEnut Ltd.	KR			PAC CLS		
Total runs:		PAC: **48**	PAC: **25** **CLS: 27**	PAC: **30** **CLS: 25**	PSG: **31**	PSG: **18**

The gray shading are a means to distinguish the consecutive table lines and has no other meaning

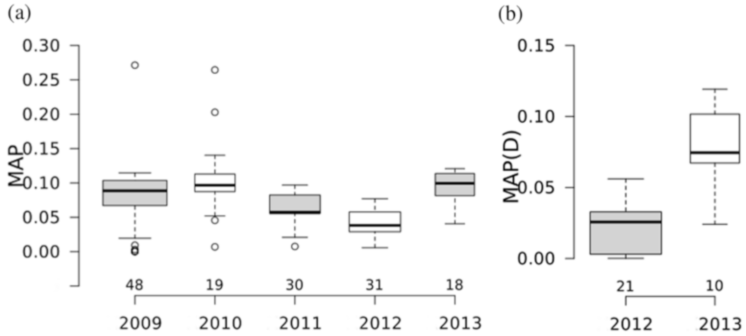

Fig. 4 Summary of MAP scores in the PAC and PSG CLEF–IP tasks. (**a**) MAP scores for the PAC tasks. (**b**) MAP(D) scores for the PSG tasks

Giachanou et al. (2013). The patent citation information stored in the document set of the collection was exploited less in 2009, with more groups using this metadata in the following years. Other very patent-specific information, like filing dates, applicant and inventor names and/or countries, was rarely used.

To give an idea of the score ranges achieved by retrieval systems participating in the Prior Art tasks, we show in Fig. 4 box plot summaries of the submitted run scores for mean average precision, MAP, and passage mean average precision, MAP(D), for each year where these tasks ran.[10] The numbers just above the years on the x-axis show the number of valid runs submitted and evaluated in the respective year. The main take away message from observing the box plots in Fig. 4 is that most IR strategies, however different in their design and methods, are equally inefficient in tackling the patent retrieval tasks. The positive outliers in these figures are, in fact, scores obtained by IR systems that integrated patent domain expertise in their design. An examples of such expertise is the query expansion with terms that do not necessarily occur in the topic patent document, but are extracted from the test collection by analysing IPC related information and/or the content of patent citations. It is also clear, from this figure, that Passage Retrieval in the patent domain, as defined by the PSG task, is an even more difficult retrieval problem.

The classification of patent documents proved to be an easier challenge than finding prior art using IR methods. This is reflected in the scores obtained by the participants' submissions. These are shown in Fig. 5 which summarises the F_1 values obtained by the experiments submitted to the CLS tasks in 2010 and 2011.

Submissions to the Classification task either used text classifiers only, like kNN or Winnow type neural networks (Derieux et al. 2010; Guyot et al. 2010; D'hondt et al. 2011), or chose a solution implementing systems similar to text retrieval

[10]Note that the scores between years cannot be directly compared, as each lab year came with a new set of test topics.

Fig. 5 CLS tasks F_1 scores

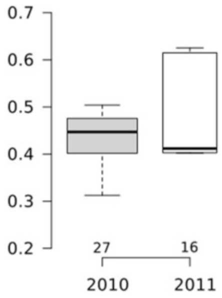

that returned the IPC codes as results, or combined classification and text retrieval (Teodoro et al. 2010; Derieux et al. 2010).

All data related to the CLEF–IP evaluation campaign (collection, topics, scripts, documentation, etc.) can be downloaded from the CLEF–IP website.[11] Detailed descriptions of the systems that participated in the CLEF–IP tasks can be found in the CLEF workshop notes available on the CLEF Initiative website[12] and on the CLEF–IP website.

6 Closing Remarks

We have presented in this chapter the development of the CLEF–IP benchmarking activity for patent text retrieval over a period of 5 years. It advanced from a simply formulated retrieval task to organizing more elaborated tasks that cover specific pieces of the Intellectual Property practitioners' daily work-flow.

At the end of the CLEF–IP evaluation campaign, it is clear to us that successful information retrieval in the patent domain involves at least well thought-out adjustments to the currently used retrieval and text mining systems to take into account the specificities of the patent domain. In general, retrieval results do not come close to the expectations of patent experts. One reason for this is that transferring the know-how of IP professionals to the IR research community is a complex undertaking. An example of such patent domain expertise which was insufficiently treated by IR researchers is language obfuscation. A method used rather often by patent applicants, language obfuscation employs vague and over-broad terms for otherwise very concrete concepts.

Even though the CLEF–IP campaign is no longer running, there is a huge potential to use the data and realistic patent search tasks resulting from the CLEF–IP campaign to develop innovative solutions in the patent information

[11] CLEF–IP: Retrieval in the Intellectual Property Domain. http://ifs.tuwien.ac.at/~clef-ip/.

[12] The CLEF Initiative (Conference and Labs of the Evaluation Forum, formerly known as Cross-Language Evaluation Forum). http://www.clef-initiative.eu/.

retrieval domain. The CLEF–IP tasks described in this chapter are focused on text-oriented information retrieval. There remains however extensive work to be done on improving the use of non-textual patent data in patent search. Early steps in this direction were done by the organisation of additional tasks, where the CLEF–IP test collection was augmented with data sets pertinent to non-textual patent content: flowcharts, chemical structures, images.

Another important aspect of patent retrieval, which was not addressed by the CLEF–IP campaign, is that information search is session based: the final list of relevant documents is the result of several search queries, possibly building on each other. Both these research directions need sustained support from the IP community.

Undertakings like TREC–CHEM, CLEF–IP, NTCIR workshop series are ambitious from at least two points of view. On one side, by interfacing with patent practitioners, these evaluation activities can be used to showcase advances in IR methods, methods that should easily be adaptable to the IP domain, and facilitate their daily need for specific information needs, allowing them to explore the patent data in novel ways. On the other side, such evaluation campaigns repeatedly bring to the attention of academic IR researchers the fact that there exists a large body of technological know-how, namely patent databases. The CLEF–IP benchmark contributed to creating a picture of the search result quality the IR methods deliver when faced with an information need like the one represented by the patent novelty search (i.e. finding relevant patents for a given patent application). The availability of patent-based test collections has triggered research in various IR areas, an inventory of the latest IP-relevant studies being also presented in Lupu and Hanbury (2013) and Lupu et al. (2017).

Acknowledgements We give our thanks to the following people:

- the advisory board members which helped shape the evaluation lab in its early years: Gianni Amati, Atsushi Fujii, Makoto Iwayama, Kalervo Järvelin, Noriko Kando, Javier Pose Rodríguez, Mark Sanderson, Henk Thomas, Anthony Trippe, Christa Womser-Hacker
- the numerous patent experts that helped us understand many of the patent system's subtilities
- previous CLEF–IP organizers and co-organizers: Giovanna Roda, John Tait, Veronika Zenz, Mihai Lupu, Igor Filippov, Walid Magdy, Alan P. Sexton
- the participating teams who submitted such a variety of solutions to the proposed tasks (see CLEF–IP workshop notes).

CLEF–IP was supported along the years by Matrixware GmbH, Vienna and Information Retrieval Facility, Vienna as first data and infrastructure provider, by the PROMISE, EU Network of Excellence (FP7-258191), and FFG FIT-IT Impex project (No. 825846).

References

Adams S (2000) Using the International Patent Classification in an online environment. World Patent Information 22(4):291–300

Adams SR (2011) Information sources in patents, 3rd edn. K.G. Saur, Munich

Alberts D, Yang CB, Fobare-DePonio D, Koubek K, Robins S, Rodgers M, Simmons E, DeMarco D (2011) Introduction to patent searching. In: Lupu M, Mayer K, Tait J, Trippe AJ (eds) Current challenges in patent information retrieval, the information retrieval series, vol 29. Springer, Berlin, pp 3–43. http://dx.doi.org/10.1007/978-3-642-19231-9_1

Anthon C (1841) A classical dictionary: containing an account of the principal proper names mentioned in ancient authors, and intended to elucidate all the important points connected with the geography, history, biography, mythology, and fine arts of the Greeks and Romans together with an account of coins, weights, and measures, with tabular values of the same. Harper & Bros

Attar R, Fraenkel AS (1977) Local feedback in full-text retrieval systems. J ACM 24(3):397–417. http://doi.acm.org/10.1145/322017.322021

Becks D, Womser-Hacker C, Mandl T, Kölle R (2009) Patent retrieval experiments in the context of the CLEF IP Track 2009. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Becks D, Mandl T, Womser-Hacker C (2010) Phrases or terms? The impact of different query types. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

Correa S, Buscaldi D, Rosso P (2009) NLEL-MAAT at CLEF-IP. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Derieux F, Bobeica M, Pois D, Raysz JP (2010) Combining semantics and statistics for patent classification. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

D'hondt E, Verberne S, Alink W, Cornacchia R (2011) Combining document representations for prior-art retrieval. In: Petras V, Forner P, Clough P, Ferro N (eds) (2011) CLEF 2011 Working Notes, CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1177/

Eiselt A, Oberreuter G (2013) The simpler the better - Retrieval Moderl comparison for Prior-Art Search in Patents at CLEF-IP 2013. In: Forner P, Navigli R, Tufis D, Ferro N (eds) (2013) CLEF 2013 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1179/

EPO (1973) European Patent Convention (EPC), Implementing Regulations, Examination Procedure. http://www.epo.org/law-practice/legal-texts/html/epc/2016/e/r71.html. Accessed Dec 2018

EPO (2018) Guidelines for Examination in the European Patent Office. Directorate Patent Law 5.2.1. http://www.epo.org/law-practice/legal-texts/guidelines.html. Accessed Dec 2018

Frietsch R, Schmoch U, Looy B, Walsh P, Devroede R, Du Plessis M, Jung T, Meng Y, Neuhäusler P, Peeters B, Schubert T (2010) The value and indicator function of patents. Expertenkommission Forschung und Innovation (EFI) Studien zum deutschen Innovationssystem 15(15-2010)

Galasso A, Schankerman M (2013) Do patents help or hinder innovation? World Economic Forum. https://www.weforum.org/agenda/2013/05/do-patents-help-or-hinder-innovation. Accessed Dec 2018

Giachanou A, Salampasis M, Satratzemi M, Samaras N (2013) Report on the CLEF-IP 2013 experiments: multilayer collection selection on topically organized patents. In: Forner P, Navigli R, Tufis D, Ferro N (eds) (2013) CLEF 2013 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1179/

Gobeill J, Ruch P (2012) BiTeM site report for the Claims to Passage task in CLEF-IP 2012. In: Forner P, Karlgren J, Womser-Hacker C, Ferro N (eds) (2012) CLEF 2012 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1178/

Gobeill J, Theodoro D, Ruch P (2009) Exploring a wide range of simple pre and post processing strategies for patent searching in CLEF IP 2009. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Graf E, Azzopardi L (2008) A methodology for building a patent test collection for prior art search. In: Proceedings of the second international workshop on evaluating information access (EVIA)

Graf E, Azzopardi L, van Rijsbergen K (2009) Automatically generating queries for prior art search. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Guyot J, Benzineb K, Falquet G (2010) myClass: a mature tool for patent classification. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

Hall BH (2017) Patents. Palgrave, Macmillan, pp 1–9. https://doi.org/10.1057/978-1-349-95121-5_1393-2

Hanbury A, Zenz V, Berger H (2010) 1st international workshop on advances in patent information retrieval (AsPIRe'10). SIGIR Forum 44(1):19–22. http://doi.acm.org/10.1145/1842890.1842893

Hunt D, Nguyen L, Rodgers M (2007) Patent searching: tools & techniques. Wiley, New York

Iwayama M, Fujii A, Kando N, Marukawa Y (2003) An empirical study on retrieval models for different document genres: patents and newspaper articles. In: Proceedings of the 26th international ACM SIGIR conference on research and development in information retrieval. ACM, New York, SIGIR '03, pp 251–258

Järvelin K, Kekäläinen J (2002) Cumulated gain-based evaluation of IR techniques. ACM Trans Inf Syst 20(4):422–446. http://doi.acm.org/10.1145/582415.582418

Kamps J, Pehcevski J, Kazai G, Lalmas M, Robertson S (2008) INEX 2007 evaluation measures. In: Focused access to XML documents, 6th international workshop of the initiative for the evaluation of XML retrieval, INEX 2007, Dagstuhl Castle, Germany, December 17–19, 2007. Selected papers. Lecture notes in computer science, vol 4862. Springer, Berlin, pp 24–33

Kando N, Leong MK (2000) Workshop on patent retrieval (SIGIR 2000 Workshop Report). SIGIR Forum 34(1):28–30

Kumagai KI (2005) History of Japanese Industrial Property System. http://www.jpo.go.jp/torikumi_e/kokusai_e/training/textbook/pdf/History_of_Japanese_Industrial_Property_System(2005).pdf. Accessed Feb 2018

Lipani A, Palotti J, Lupu M, Piroi F, Zuccon G, Hanbury A (2017) Fixed-cost pooling strategies based on IR evaluation measures. In: Advances in information retrieval - 39th European conference on IR research, ECIR 2017, Aberdeen, April 8–13, 2017, Proceedings, pp 357–368. https://doi.org/10.1007/978-3-319-56608-5_28

Lopez P, Romary L (2009) Multiple retrieval models and regression models for prior art search. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Lopez P, Romary L (2010) Experiments with citation mining and key-term extraction for prior art search. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

Lupu M, Hanbury A (2013) Patent retrieval. Found Trends IR 7:1–97

Lupu M, Huang J, Zhu J, Tait J (2009) TREC-CHEM: large scale chemical information retrieval evaluation at TREC. SIGIR Forum 43(2):63–70

Lupu M, Mayer K, Kando N, Trippe A (2017) Current challenges in patent information retrieval, 2nd edn. Springer, Berlin. https://doi.org/10.1007/978-3-662-53817-3

Magdy W, Jones G (2010a) PRES: a score metric for evaluating recall-oriented information retrieval applications. In: Proceedings of the 33rd international ACM SIGIR conference on research and development in information retrieval, ACM, New York, SIGIR '10, pp 611–618. http://doi.acm.org/10.1145/1835449.1835551

Magdy W, Jones GJF (2010b) Examining the robustness of evaluation metrics for patent retrieval with incomplete relevance judgements. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

Magdy W, Leveling J, Jones G (2009) DCU at CLEF-IP 2009: exploring standard IR techniques on patent retrieval. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Mahdabi P, Andersson L, Hanbury A, Crestani F (2011) Report on the CLEF-IP 2011 experiments: exploring patent summarization. In: Petras V, Forner P, Clough P, Ferro N (eds) (2011) CLEF 2011 Working Notes, CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1177/CLEF2011wn-CLEF-IP-MahdabiEt2011.pdf

May C (2010) The Venetian moment: new technologies, legal innovation and the institutional origins of intellectual property. Prometheus 20:159–179. http://www.tandfonline.com/doi/full/10.1080/08109020210138979

McDonnell P (1969) Technical information management in the U.S. patent office. J Chem Doc 9(5):220–224

Mossoff A (2007) Who cares What Thomas Jefferson thought about patents - reevaluating the patent privilege in historical context. Cornell Law Rev 92(5):953. https://scholarship.law.cornell.edu/clr/vol92/iss5/2

Osborn M, Strzalkowski T, Marinescu M (1997) Evaluating document retrieval in patent database: a preliminary report. In: Proceedings of the 6th international conference on information and knowledge management. ACM, New York, pp 216–221. http://doi.acm.org/10.1145/266714.266899

PCT (1970) Patent Cooperation Treaty. http://www.wipo.int/pct/en/treaty/about.html. Accessed Aug 2015

Pfaller W (2013a) Bergrecht, Monopole, Privilegien. http://www.wolfgang-pfaller.de/berg.htm. Accessed Feb 2018

Pfaller W (2013b) Schon die alten Griechen .... http://www.wolfgang-pfaller.de/sybaris.htm. Accessed Feb 2018

Piroi F, Zenz V (2011) Evaluating information retrieval in the intellectual property domain: the CLEF-IP campaign. In: Lupu M, Mayer K, Tait J, Trippe AJ (eds) Current challenges in patent information retrieval, the information retrieval series, vol 29. Springer, Berlin, pp 87–108. http://dx.doi.org/10.1007/978-3-642-19231-9_4

Piroi F, Lupu M, Hanbury A, Zenz V (2011) CLEF-IP 2011: retrieval in the intellectual property domain. In: Petras V, Forner P, Clough P, Ferro N (eds) (2011) CLEF 2011 Working Notes, CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1177/

Piroi F, Lupu M, Hanbury A, Sexton AP, Magdy W, Filippov IV (2012) CLEF-IP 2012: retrieval experiments in the intellectual property domain. In: Forner P, Karlgren J, Womser-Hacker C, Ferro N (eds) (2012) CLEF 2012 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1178/

Piroi F, Lupu M, Hanbury A (2013) Overview of CLEF-IP 2013 Lab - information retrieval in the patent domain. In: Forner P, Müller H, Paredes R, Rosso P, Stein B (eds) Information access evaluation meets multilinguality, multimodality, and visualization. Proceedings of the fourth international conference of the CLEF initiative (CLEF 2013). Lecture Notes in Computer Science (LNCS), vol 8138. Springer, Heidelberg

Rich GS (1993) Are Letters Patent Grants of Monopoly? Western New England Law Review 15. https://core.ac.uk/display/76563380

Roda G, Tait J, Piroi F, Zenz V (2010) CLEF-IP 2009: retrieval experiments in the intellectual property domain. In: Peters C, Nunzio GD, Kurimo M, Mostefa D, Penas A, Roda G (eds) Multilingual information access evaluation I. Text retrieval experiments 10th workshop of the cross-language evaluation forum, CLEF 2009, vol 6241. Springer, Berlin, pp 385–409

Schneller I (2002) Japanese File Index Classification and F-terms. World Patent Inf 24(3):197–201

Skolnik H (1977) Historical aspects of patent systems. J Chem Inf Comput Sci 17(3):119–121. https://pubs.acs.org/doi/abs/10.1021/ci60011a002

Spark-Jones K, Van Rijsbergen C (1975) Report on the need for and provision of an 'ideal' information retrieval test collection. 5266, Computer Laboratory, Univ. Cambridge. http://sigir.org/files/museum/pub-14/pub_14.pdf

Szarvas G, Herbert B, Ggurevych I (2009) Prior art search using international patent classification codes and all-claims-queries. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Tait J, Harris C, Lupu M (eds) (2010) PaIR '10: proceedings of the 3rd international workshop on patent information retrieval. ACM, New York

Teodoro D, Gobeill J, Pasche E, Vishnyakova D, Ruch P, Lovis C (2010) Automatic prior art searching and patent encoding at CLEF-IP'10. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

Toucedo C, Losada D (2009) University of Santiago de Compostella at CLEF-IP'09. In: Borri F, Nardi A, Peters C, Ferro N (eds) (2009) CLEF 2009 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1175/

Verberne S, D'hondt E (2011) Patent classification experiments with the Linguistic Classification System LCS in CLEF-IP 2011. In: Petras V, Forner P, Clough P, Ferro N (eds) (2011) CLEF 2011 Working Notes, CEUR Workshop Proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1177/

Voorhees EM, Harman DK (2005) TREC: experiment and evaluation in information retrieval (digital libraries and electronic publishing). The MIT Press, Cambridge

Wanagiri M, Adriani M (2010) Prior art retrieval using various patent document fields contents. In: Braschler M, Harman DK, Pianta E, Ferro N (eds) (2010) CLEF 2010 working notes. CEUR workshop proceedings (CEUR-WS.org), ISSN 1613-0073. http://ceur-ws.org/Vol-1176/

WIPO (2013) Handbook on industrial property information and documentation. Part 8: terms and abbreviations concerning industrial property information and documentation. http://www.wipo.int/standards/en/part_08.html. Accessed Jan 2019