Overview of the Patent Retrieval Task at the NTCIR-6 Workshop

Atsushi Fujii*, Makoto Iwayama†, Noriko Kando‡

*Graduate School of Library, Information and Media Studies
University of Tsukuba
1-2 Kasuga, Tsukuba, 305-8550, Japan
fujii@slis.tsukuba.ac.jp

†Hitachi, Ltd.
1-280 Higashi-Kougakubo, Kokubunji, 185-8601, Japan
makoto.iwayama.nw@hitachi.com
/ Tokyo Institute of Technology

‡ National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-ku, 101-8430, Japan
kando@nii.ac.jp

Abstract

In the Sixth NTCIR Workshop, we organized the Patent Retrieval Task and performed three subtasks; Japanese Retrieval, English Retrieval, and Classification. This paper describes the Japanese Retrieval Subtask and English Retrieval Subtask, both of which were intended for patent-to-patent invalidity search task. We report the evaluation results of the groups participating in those subtasks.

Keywords: *Patent retrieval, Invalidity search, Patent classification, NTCIR*

1 Introduction

Processes of patent retrieval differ significantly, depending on the purpose of retrieval. One process is the "technology survey", in which patents related to a specific technology, e.g., "blue light-emitting diode", are searched. This process is similar to ad hoc retrieval tasks targeting non-patent documents.

Another process is the "invalidity search", in which prior arts related to a patent application are searched. Away from academic research, invalidity searches are performed by examiners in government patent offices and searchers in the intellectual property divisions of private companies.

In the Third NTCIR Workshop (NTCIR-3), the authors of this paper organized the Patent Retrieval Task [8]. In NTCIR-3, we performed the technology survey task, in which patents were used as technical publications.

Given a success in NTCIR-3, we also performed the Patent Retrieval Task in the Fourth NTCIR Workshop (NTCIR-4) focusing on the invalidity search and patent map generation subtasks [1, 2, 3].

In NTCIR-5, we performed three subtasks; Document Retrieval Subtask, Passage Retrieval Subtask, and Classification Subtask [4, 5, 6]. The Document Retrieval Subtask was the same as the invalidity search in NTCIR-4. However, the numbers of search topics and target documents were increased. The purpose of the Passage Retrieval Subtask was to sort passages in a document according to the relevance to a topic. The purpose of the Classification Subtask was to categorize patent documents based on a multi-dimensional classification system called "F-term (File Forming Term)".

In NTCIR-6, we performed the following subtasks:

- Japanese Retrieval Subtask

  The purpose was the same as the Document Retrieval Subtask in NTCIR-5. However, we introduced a new relevance degree.

- English Retrieval Subtask

  The purpose was the invalidity search as in the Japanese Retrieval Subtask. However, we used patent grant data published by the U.S. Patent & Trademark Office (USPTO) as target documents.

- Classification Subtask

  The purpose was the same as the Classification Subtask in NTCIR-5. However, the number of topics was increased. We also explored evaluation methods using the hierarchy of the F-term.

After the NTCIR-6 Workshop meeting, the test collections for these subtasks will be available to the public for research purposes[1].

[1] http://if-lab.slis.tsukuba.ac.jp/fujii/ntc6pat/cfp-en.html

This paper describes the Japanese Retrieval Subtask and the English Retrieval Subtask. The Classification Subtask is described by Iwayama et al. [7].

2 Japanese Retrieval Subtask

2.1 Overview

The purpose of the invalidity search is to find one or more patents that can invalidate the demand in an existing claim. This is a patent-to-patent associative retrieval task. Away from academic research, invalidity searches are usually performed by examiners in a government patent office and searchers of the intellectual property division in private companies.

The Japanese Retrieval Subtask was performed as follows. First, the task organizers (i.e., the authors of this paper) provided each participating group with a document collection and search topics.

Second, each group submitted retrieval results obtained by the topics. Each group was allowed to submit one or more retrieval results. In a single retrieval result, up to the top 1000 retrieved documents must be sorted by the relevance score.

Finally, the organizers evaluated the submitted results using relevant documents. The evaluation results were sent to each group, who was also encouraged to analyze the results of their methods and report the obtained knowledge at the workshop meeting.

The difference from the Document Retrieval Subtask in NTCIR-5 is three fold.

First, the definition of the relevance was changed. During NTCIR-5, we identified that the retrieval task was difficult when a topic patent application and the citation used to reject the application do not share the same IPC codes [4]. Thus, we defined relevant documents and partially relevant document according to the IPC. See Section 2.5 for details.

Second, to compare the evaluation results across the past test collections, the participating groups were requested to submit retrieval results for the four different topic sets; NTCIR-4, NTCIR-5, the dry run of NTCIR-6, and the formal run of NTCIR-6.

Finally, in each retrieval result (i.e., run), the number of top documents to be examined must be specified. This number is 1000 or less than 1000. This rule is intended to improve the utility of best-match retrieval systems. Unlike exact-match retrieval systems, which identify the exact number of documents to be examined, in best-match retrieval systems it is not obvious how many top documents must be examined. This is one reason why professional patent searchers still use exact-match retrieval systems. However, because most of the participating groups did not address this problem, we did not apply this rule in the evaluation of the formal run.

2.2 Document Collection

In NTCIR-4, the document collection consisted of five years of unexamined Japanese patent applications published in 1993–1997. In NTCIR-5, the document collection consisted of ten years of unexamined Japanese patent applications published in 1993–2002. The number of documents in the collection was 3,496,252. In NTCIR-6, we used the same document collection as NTCIR-5.

To standardize the format of the documents, the organizers provided an official tool, which inserts SGML-style tags into each document. Table 1 shows the tags inserted by that tool. Although <PASSAGE> and <PNUM> were intended for the Passage Retrieval Subtask in NTCIR-5, the participating groups were allowed to use passage information for document retrieval purposes in NTCIR-6.

Table 1. Tags for Japanese patent applications.

<DOC>	document
<DOCNO>	document identifier
<TEXT>	text body
<PASSAGE>	passage
<PNUM>	passage identifier

2.3 Search Topics

Each search topic is a Japanese patent application rejected by the Japanese Patent Office (JPO). For each topic, one or more citations (i.e., prior art) were identified by examiners of the JPO to invalidate the demand in the topic patent.

To increase the number of topics, we decreased the cost required for producing search topics and relevance judgements as much as possible. We automatically extracted patent applications rejected by the JPO and the citations used for the rejection. For this purpose, we used the citation information in the "*Seirihyoujunka* (Standardized)" Data, which had been produced from the master database in the JPO. We used only the citations as relevant or partially relevant documents and did not perform relevance judgement by human assessors.

An application used as a search topic must satisfy all of the following four criteria.

- the application was not used as a search topic in NTCIR-4 and NTCIR-5.
- the texts of both the application and the corresponding citation are included in the document collection.

- the citation had been published before the application was filed because to invalidate the invention in a topic patent, relevant documents must be prior art.
- the application does not claim the priority; otherwise it is difficult to identify the filing date of the application automatically.

From applications satisfying the above criteria, we selected 1685 applications as search topics. Although the number of topics was determined with no particular reason, we intended to produce more than one thousand topics. To produce topics with a large number of relevant documents, we selected applications for which at least five citations are provided.

Because candidates of relevant documents for an application can be limited by the filing date of the application, for old applications the retrieval of relevant documents can be a trivial task. Thus, we selected search topics from recently filed applications.

The citation information we used did not include the information as to which claim was the target of the rejection. From each application, we systematically extracted the first claim, which is usually the target, as a topic.

Each search topic file includes a number of additional SGML-style tags. The claim used as the target of invalidation is specified by <CLAIM>. The date of filing is specified by <FDATE> and only the patents published before this date can potentially be relevant.

Table 2 shows the tags inserted into each topic. Because each topic is a Japanese patent application, the tags in Table 1 are also inserted in each topic.

Table 2. Tags for Japanese topics.

<TOPIC>	topic
<NUM>	topic identifier
<LANG>	topic language
<PURPOSE>	purpose of search
<FDATE>	date of filing
<CLAIM>	target claim

For the dry run, instead of the "*Seirihyoujunka* (Standardized)" Data, we produced the topics from a collection of search reports produced by professional patent search intermediaries. The source data set was reference data for examiners at the JPO. We used 1000 reports produced in 2001–2003, from which we produced 349 topics.

2.4 Submissions

Each participating group was allowed to submit one or more retrieval results, in which at least one result must be obtained using only the <CLAIM> and <FDATE> fields. For the remaining results, any information in a topic file, such as IPC codes, can be used.

The participating groups were requested to submit retrieval results for the following topic sets:

- NTCIR-4: 34 topics
- NTCIR-5: 1189 topics
- Search report data: 349 topics for the dry run
- NTCIR-6: 1685 topics

In NTCIR-4, although a small number of topics were produced, for each topic professional searchers performed relevance judgement and collected as many relevant documents as possible. In NTCIR-5 and NTCIR-6, although a larger number of topics were produced, for each topic only the citations provided by examiners in the JPO, which are not exhaustively listed, were used as the relevant documents. In the Search report data, the number of topics and the exhaustiveness for the relevant documents are between those for NTCIR-4 and NTCIR-5/6.

2.5 Evaluation Method

For the Document Retrieval Subtask in NTCIR-5, the relevance degree of the citation for to a topic was determined based on the following two ranks:

- the citation used to reject an application was regarded as a "relevant document (A)" because the decision was made confidently,
- a citation used to reject an application with another citation was regarded as a "partially relevant document (B)", because each citation is partially related to the claim in the application.

We call this definition "Def0".

In NTCIR-6, we used an alternative definition for the relevance degree. If a topic patent and its relevant document are assigned to the same IPC code, the document can usually be retrieved with a high accuracy. We define the following two sets for explanation purposes.

- TS: a set of IPC subclasses assigned to topic T
- DS: a set of IPC subclasses assigned to relevant document D

The degree of the relevance of D was classified into the following three ranks.

- H: the intersection of TS and DS is empty.
- A: the intersection of TS and DS is not empty, but TS and DS are not identical.
- B: TS and DS are identical.

We call this definition "Def1".

We used Mean Average Precision (MAP), which has commonly been used in past IR literature, to evaluate the submitted runs for the Japanese Retrieval Subtask. In Section 4, we show the MAP values for Def0 and Def1, respectively.

2.6 Policy for Resource Usage

Because the citations provided by the examiners of the JPO are available to the public, participating groups can obtain relevant documents for the topics before the formal run.

However, participating groups were not allowed to use the citations corresponding to a topic application for training purposes. In addition, because we used search topics in the past test collections for Patent Retrieval Tasks, the relevance judgements for these test collections must not be used for training purposes.

These rules were not applied if their system had already been trained by a large number of citations and it was not easy to remove information, such as statistics, related to specific citations.

Except for the above cases, participating groups were allowed to use any resources for the task according to the "reasonable use policy".

3 English Retrieval Subtask

3.1 Overview

In the English Retrieval Subtask, we used patent grant data published from the USPTO as target documents. Each topic is also a USPTO patent and the purpose is to invalidate the demand in the topic. In each patent, the applicant and the examiner list one or more citations as prior arts. We used the citations in a topic patent as relevant documents for the topic. As in the Japanese Retrieval Subtask, we used the "reasonable use policy" for resource usage.

3.2 Document Collection

The document collection consists of eight years of patent grant data published in 1993–2000. The number of documents in the collection is 981,948. While in the Japanese Retrieval Subtask we used patent applications as target documents, in the English Retrieval Subtask we used only patents that have been granted by the USPTO.

To standardize the format of the documents, the organizers provided an official tool, which inserts SGML-style tags into each document. Table 3 shows the tags inserted by that tool. Because the format of the source data was more complicated than that for the Japanese patent applications, we inserted a large number of tags to enhance the readability of the USPTO patent data.

Table 3. Tags for English patents.

<DOC>	document
<DOCNO>	document identifier
<APP-NO>	application number
<APP-DATE>	application date
<PUB-NO>	publication number
<PUB-TYPE>	publication type
<PAT-NO>	patent number
<PAT-TYPE>	patent type
<PUB-DATE>	publication date
<PRI-IPC>	primary IPC
<IPC-VER>	IPC version
<PRI-USPC>	primary USPC
<PRIORITY>	priority information
<CITATION>	citation(s)
<INVENTOR>	inventor(s)
<ASSIGNEE>	assignee(s)
<TITLE>	title
<ABST>	abstract
<SPEC>	specification
<CLAIM>	claim(s)

3.3 Search Topics

Each search topic is a patent grant published in 2000–2001. We selected patents that satisfy the following criteria:

- at least 20 citations are listed.
- at least 90% of the citations are included in the target document collection.

As a result, we collected 3221 topics, in which we used 1000 topics for training purposes and 2221 topics for the formal run. The format of the topic is the same as that for the target documents. In addition to the tags in Table 3, we inserted the tags in Table 2 except for <FDATE> into each topic.

3.4 Submissions

Each participating group was allowed to submit one or more retrieval results, in which at least one result must be obtained using only the <CLAIM> field. For the remaining results, any information in a topic file can be used.

3.5 Evaluation Method

The relevant patents for a search topic is one or more citations provided by the applicant and the ex-

aminer in the topic patent.

If a topic patent and its relevant document are assigned to the same IPC class, the document can usually be retrieved with a high accuracy. Thus, the degree of the relevance of each citation was classified into the following two ranks.

- A: The IPC subclasses assigned to the topic patent and the target document are not identical.
- B: The IPC subclasses assigned to the topic patent and the target document are identical.

Unlike Japanese patent applications, only a single IPC code is assigned to each USPTO patent. Thus, we did not consider TS and DS, which was used in Section 2.5 for explanation purposes.

We used Mean Average Precision (MAP) to evaluate the submitted runs for the English Retrieval Subtask.

4 Evaluation

4.1 Overview

For the Japanese Retrieval Subtask and the English Retrieval Subtask, we used MAP as the evaluation measure. To calculate MAP for each submitted run, the organizers produced a Perl program that is compatible with trec_eval. This Perl program was also used for the evaluation in the Classification Subtask. The value of MAP can potentially be different depending on the version of trec_eval.

4.2 Japanese Retrieval Subtask

Table 4 shows the submitted runs for the Japanese Retrieval Subtask, which also shows the submission type (Mandatory or Optional) for each run. From five groups, 22 runs were submitted.

Tables 5–9 show the MAP values of the submitted runs for different combinations of different conditions; relevance and relevance degree. The condition can be identified with the caption of each table. “Def0” and “Def1” correspond to the definitions described in Section 2.5. “Rel” denotes the relevance degree.

In Tables 5–9, the column “NTCX” denotes MAP for NTCIR-X topics, for which X is “4”, “5”, or “6”. the column “SR” denotes MAP for the search reports. In the column “Run ID”, run IDs are sorted according the MAP values in “Total”.

For “Def0”, “H” relevant documents do not exist. By definition, “A” relevant documents do not exist for “Def0” in “NTC6”.

Looking at Tables 5–9, the relative superiority between groups in MAP was almost the same, irrespective of the condition. However, the relative superiority between runs of the same group in MAP was salient depending on the condition.

4.3 English Retrieval Subtask

Table 10 shows the submitted runs for the English Retrieval Subtask, which also shows the submission type (Mandatory or Optional) for each run. From five groups, 11 runs were submitted. However, JSPAT0–3, which were received after the submission deadline, are not regarded as the official results.

Table 10 also shows the MAP values of the submitted runs for different relevance degrees. While in “A” only “A” relevant documents were used as the correct answers, in “AB” both “A” and “B” relevant documents were used as the correct answers.

Looking at Table 10, the MAP values are generally small. Because this is mainly due to the selection of the search topics, we need to explore the process of producing search topics.

5 Conclusion

In the Sixth NTCIR Workshop, we organized the Patent Retrieval Task and performed three subtasks. This paper described the Japanese Retrieval Subtask and English Retrieval Subtask. Both subtasks were intended for the patent-to-patent invalidity search.

During the evaluation in the Japanese Retrieval Subtask, we compared the results obtained by the topics with exhaustively listed relevance judgement and the topics with limited relevance judgement. However, the exhaustiveness for relevance judgement did not affect the relative superiority between participating groups in the retrieval accuracy.

Table 4. Submitted runs for Japanese Retrieval Subtask.

Run ID	Type
AFLAB1	Mandatory
BETA6-1	Mandatory
hcu1	Mandatory
hcu2	Mandatory
HTC01	Mandatory
HTC02	Mandatory
HTC03	Mandatory
HTC04	Optional
HTC05	Optional
HTC06	Optional
HTC07	Optional
HTC08	Optional
HTC09	Optional
HTC10	Optional
JSPAT0	Mandatory
JSPAT1	Mandatory
JSPAT2	Mandatory
JSPAT3	Mandatory
JSPAT4	Mandatory
JSPAT5	Mandatory
JSPAT6	Mandatory
JSPAT7	Mandatory

Table 5. MAP for Japanese Retrieval Subtask: Def0, Rel=A

Run ID	Total	NTC4	NTC5	SR	NTC6
HTC09	23.23	28.28	23.91	18.32	—
HTC10	23.22	28.34	23.91	18.25	—
HTC08	21.79	28.12	22.43	16.80	—
HTC07	21.72	28.56	22.24	17.19	—
HTC06	21.00	24.22	21.65	16.74	—
HTC05	19.90	30.14	20.11	16.12	—
HTC04	19.83	34.09	19.86	15.91	—
HTC02	19.55	29.93	19.67	16.20	—
HTC03	19.14	29.73	19.17	16.25	—
AFLAB1	18.38	21.37	19.16	13.50	—
HTC01	18.37	27.79	18.52	15.10	—
JSPAT3	10.47	12.36	10.92	7.58	—
JSPAT7	10.44	12.24	10.90	7.57	—
JSPAT0	10.39	12.11	10.83	7.60	—
JSPAT4	10.36	11.98	10.81	7.59	—
JSPAT1	10.21	11.80	10.63	7.60	—
JSPAT5	10.19	11.69	10.61	7.59	—
hcu1	9.78	20.52	9.69	7.39	—
hcu2	9.62	19.38	9.54	7.47	—
JSPAT2	9.56	13.92	9.99	6.10	—
JSPAT6	9.54	13.91	9.97	6.07	—
BETA6-1	6.92	14.39	6.94	4.86	—

Table 6. MAP for Japanese Retrieval Subtask: Def0, Rel=AB

Run ID	Total	NTC4	NTC5	SR	NTC6
HTC10	15.41	26.36	19.73	18.45	11.51
HTC09	15.29	26.27	19.67	18.34	11.34
HTC08	14.45	25.31	18.56	17.17	10.77
HTC07	14.41	24.85	18.47	17.44	10.71
HTC06	13.96	22.66	17.95	16.63	10.41
HTC05	12.96	24.18	16.60	15.99	9.53
HTC04	12.12	25.96	16.30	14.38	8.43
HTC02	11.72	24.09	15.77	14.05	8.12
HTC03	11.66	24.11	15.48	14.43	8.15
AFLAB1	11.46	16.15	15.39	13.27	8.21
HTC01	11.13	21.23	15.03	12.76	7.83
hcu1	6.49	15.67	8.22	7.52	4.88
hcu2	6.37	15.30	8.06	7.47	4.77
JSPAT3	6.26	12.06	8.77	6.47	4.31
JSPAT0	6.23	10.95	8.54	6.65	4.41
JSPAT7	6.22	11.12	8.75	6.45	4.28
JSPAT1	6.21	10.73	8.55	6.66	4.38
JSPAT4	6.18	10.08	8.50	6.63	4.37
JSPAT5	6.17	9.87	8.52	6.64	4.34
JSPAT2	5.61	11.51	8.06	5.30	3.81
JSPAT6	5.58	11.17	8.04	5.27	3.79
BETA6-1	4.89	11.31	5.86	5.98	3.85

Table 7. MAP for Japanese Retrieval Subtask: Def1, Rel=H

Run ID	Total	NTC4	NTC5	SR	NTC6
HTC06	5.98	7.39	10.13	11.32	4.29
HTC08	5.83	9.17	9.09	12.58	4.25
HTC07	5.74	9.23	9.03	11.88	4.20
HTC05	4.98	6.45	7.87	10.16	3.67
HTC04	4.54	10.32	8.25	8.46	2.99
HTC10	4.34	8.40	6.32	6.33	3.48
HTC03	4.22	9.35	7.25	9.01	2.82
HTC02	3.93	8.69	6.99	5.51	2.77
HTC01	3.78	7.37	6.15	6.65	2.76
AFLAB1	3.66	10.03	5.36	7.45	2.71
HTC09	3.01	8.15	4.71	3.69	2.31
hcu2	2.56	3.17	4.41	6.75	1.67
hcu1	2.54	3.19	4.34	6.23	1.70
JSPAT1	1.96	5.49	2.97	2.54	1.52
JSPAT5	1.95	5.45	2.96	2.48	1.51
JSPAT3	1.95	6.06	2.44	1.22	1.76
JSPAT7	1.94	6.04	2.43	1.16	1.75
JSPAT0	1.87	5.66	2.64	2.50	1.51
JSPAT4	1.86	5.63	2.63	2.43	1.49
JSPAT2	1.76	13.08	2.18	0.95	1.42
JSPAT6	1.75	13.06	2.17	0.91	1.41
BETA6-1	0.14	1.37	0.13	0.00	0.13

Table 8. MAP for Japanese Retrieval Subtask: Def1, Rel=HA

Run ID	Total	NTC4	NTC5	SR	NTC6
HTC10	12.63	19.84	16.66	17.33	9.65
HTC09	12.46	19.70	16.52	17.14	9.46
HTC08	11.87	18.12	15.74	16.02	9.06
HTC07	11.84	17.39	15.82	16.12	8.97
HTC06	11.50	15.32	15.41	15.61	8.73
HTC05	10.60	17.34	14.39	14.24	7.91
HTC04	9.88	19.50	13.92	13.12	7.09
HTC02	9.27	17.82	12.88	12.65	6.68
AFLAB1	9.25	12.59	13.24	11.77	6.69
HTC03	9.19	18.47	12.52	13.01	6.68
HTC01	9.02	15.13	12.54	11.64	6.65
hcu1	5.33	10.77	7.01	6.58	4.16
hcu2	5.26	10.35	7.03	6.44	4.06
JSPAT3	4.89	6.56	6.96	4.90	3.78
JSPAT7	4.87	5.71	6.95	4.88	3.76
JSPAT1	4.81	4.92	6.63	5.45	3.75
JSPAT0	4.79	5.02	6.62	5.40	3.73
JSPAT5	4.78	4.16	6.60	5.43	3.73
JSPAT4	4.76	4.23	6.59	5.39	3.71
JSPAT2	4.36	8.32	6.36	4.32	3.25
JSPAT6	4.34	8.22	6.35	4.30	3.23
BETA6-1	3.61	7.51	3.86	4.72	3.24

Table 9. MAP for Japanese Retrieval Subtask: Def1, Rel=HAB

Run ID	Total	NTC4	NTC5	SR	NTC6
HTC10	15.41	26.36	19.73	18.45	11.51
HTC09	15.29	26.27	19.67	18.34	11.34
HTC08	14.45	25.31	18.56	17.17	10.77
HTC07	14.41	24.85	18.47	17.44	10.71
HTC06	13.96	22.66	17.95	16.63	10.41
HTC05	12.96	24.18	16.60	15.99	9.53
HTC04	12.12	25.96	16.30	14.38	8.43
HTC02	11.72	24.09	15.77	14.05	8.12
HTC03	11.66	24.11	15.48	14.43	8.15
AFLAB1	11.46	16.15	15.39	13.27	8.21
HTC01	11.13	21.23	15.03	12.76	7.83
hcu1	6.49	15.67	8.22	7.52	4.88
hcu2	6.37	15.30	8.06	7.47	4.77
JSPAT3	6.26	12.06	8.77	6.47	4.31
JSPAT0	6.23	10.95	8.54	6.65	4.41
JSPAT7	6.22	11.12	8.75	6.45	4.28
JSPAT1	6.21	10.73	8.55	6.66	4.38
JSPAT4	6.18	10.08	8.50	6.63	4.37
JSPAT5	6.17	9.87	8.52	6.64	4.34
JSPAT2	5.61	11.51	8.06	5.30	3.81
JSPAT6	5.58	11.17	8.04	5.27	3.79
BETA6-1	4.89	11.31	5.86	5.98	3.85

Table 10. MAP for English Retrieval Subtask.

Run ID	Type	A	AB
AFLAB1	Mandatory	3.65	7.12
AFLAB2	Optional	4.17	8.11
AFLAB3	Optional	3.81	7.48
hcu1	Mandatory	3.37	2.12
hcu2	Mandatory	3.37	2.10
JSPAT0	Mandatory	1.27	1.85
JSPAT1	Mandatory	1.26	1.84
JSPAT2	Optional	1.23	5.72
JSPAT3	Optional	1.22	4.50
KLE1	Mandatory	2.82	6.94
NTNU	Mandatory	2.30	6.94

References

[1] A. Fujii, M. Iwayama, and N. Kando. Overview of patent retrieval task at NTCIR-4. In *Proceedings of the Fourth NTCIR Workshop on Research in Information Access Technologies Information Retrieval, Question Answering and Summarization*, 2004.

[2] A. Fujii, M. Iwayama, and N. Kando. The patent retrieval task in the fourth NTCIR workshop. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 560–561, 2004.

[3] A. Fujii, M. Iwayama, and N. Kando. Test collections for patent-to-patent retrieval and patent map generation in NTCIR-4 workshop. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, pages 1643–1646, 2004.

[4] A. Fujii, M. Iwayama, and N. Kando. Overview of patent retrieval task at NTCIR-5. In *Proceedings of the Fifth NTCIR Workshop Meeting on Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-lingual Information Access*, pages 269–277, 2005.

[5] A. Fujii, M. Iwayama, and N. Kando. Test collections for patent retrieval and patent classification in the fifth NTCIR workshop. In *Proceedings of the 5th International Conference on Language Resources and Evaluation*, pages 671–674, 2006.

[6] M. Iwayama, A. Fujii, and N. Kando. Overview of classification subtask at NTCIR-5 patent retrieval task. In *Proceedings of the Fifth NTCIR Workshop Meeting on Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-lingual Information Access*, pages 278–286, 2005.

[7] M. Iwayama, A. Fujii, and N. Kando. Overview of classification subtask at NTCIR-6 patent retrieval task. In *Proceedings of the Sixth NTCIR Workshop Meeting*, 2007.

[8] M. Iwayama, A. Fujii, N. Kando, and Y. Marukawa. Evaluating patent retrieval in the third NTCIR workshop. *Information Processing & Management*, 42(1):207–221, 2006.