UvA-DARE (Digital Academic Repository)

Overview of the INEX 2008 Ad Hoc Track

Kamps, J.; Geva, S.; Trotman, A.; Woodley, A.; Koolen, M.

Publication date
2008

Published in
INEX 2008 workshop pre-proceedings

[Link to publication](#)

Citation for published version (APA):

Kamps, J., Geva, S., Trotman, A., Woodley, A., & Koolen, M. (2008). Overview of the INEX 2008 Ad Hoc Track. In S. Geva, J. Kamps, & A. Trotman (Eds.), *INEX 2008 workshop pre-proceedings* (pp. 1-28) <http://www.inex.otago.ac.nz/data/proceedings/INEX2008-preproceedings.pdf>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Overview of the INEX 2008 Ad Hoc Track

Jaap Kamps¹, Shlomo Geva², Andrew Trotman³,
Alan Woodley², and Marijn Koolen¹

¹ University of Amsterdam, Amsterdam, The Netherlands
{[kamps,m.h.a.koolen](mailto:kamps,m.h.a.koolen@uva.nl)}@uva.nl

² Queensland University of Technology, Brisbane, Australia
{[s.geva,a.woodley](mailto:s.geva,a.woodley@qut.edu.au)}@qut.edu.au

³ University of Otago, Dunedin, New Zealand
andrew@cs.otago.ac.nz

Abstract. This paper gives an overview of the INEX 2008 Ad Hoc Track. The main goals of the Ad Hoc Track were two-fold. The first goal was to investigate the value of the internal document structure (as provided by the XML mark-up) for retrieving relevant information. This is a continuation of INEX 2007 and, for this reason, the retrieval results are liberalized to arbitrary passages and measures were chosen to fairly compare systems retrieving elements, ranges of elements, and arbitrary passages. The second goal was to compare focused retrieval to article retrieval more directly than in earlier years. For this reason, standard document retrieval rankings have been derived from all runs, and evaluated with standard measures. In addition, a set of queries targeting Wikipedia have been derived from a proxy log, and the runs are also evaluated against the clicked Wikipedia pages. The INEX 2008 Ad Hoc Track featured three tasks: For the *Focused Task* a ranked-list of non-overlapping results (elements or passages) was needed. For the *Relevant in Context Task* non-overlapping results (elements or passages) were returned grouped by the article from which they came. For the *Best in Context Task* a single starting point (element start tag or passage start) for each article was needed. We discuss the results for the three tasks, and examine the relative effectiveness of element and passage retrieval. This is examined in the context of content only (CO, or Keyword) search as well as content and structure (CAS, or structured) search. Finally, we look at the ability of focused retrieval techniques to rank articles, using standard document retrieval techniques, both against the judged topics as well as against queries and clicks from a proxy log.

1 Introduction

This paper gives an overview of the INEX 2008 Ad Hoc Track. There are two main research question underlying the Ad Hoc Track. The first main research question is that of the value of the internal document structure (mark-up) for retrieving relevant information. That is, does the document structure help in identify where the relevant information is within a document? This question, first studied at INEX 2007, has attracted a lot of attention in recent years.

Trotman and Geva [11] argued that, since INEX relevance assessments are not bound to XML element boundaries, retrieval systems should also not be bound to XML element boundaries. Their implicit assumption is that a system returning passages is at least as effective as a system returning XML elements. This assumption is based on the observation that elements are of a lower granularity than passages and so all elements can be described as passages. The reverse, however is not true and only some passages can be described as elements. Huang et al. [4] implement a fixed window passage retrieval system and show that a comparable element retrieval ranking can be derived. In a similar study, Itakura and Clarke [5] show that although ranking elements based on passage-evidence is comparable, a direct estimation of the relevance of elements is superior. Finally, Kamps and Koolen [6] study the relation between the passages highlighted by the assessors and the XML structure of the collection directly, showing reasonable correspondence between the document structure and the relevant information.

Up to now, element and passage retrieval approaches could only be compared when mapping passages to elements. This may significantly affect the comparison, since the mapping is non-trivial and, of course, turns the passage retrieval approaches effectively into element retrieval approaches. To study the value of the document structure through direct comparison of element and passage retrieval approaches, the retrieval results were liberalized to arbitrary passages. Every XML element is, of course, also a passage of text. At INEX 2008, a simple passage retrieval format was introduced using file-offset-length (FOL) triplets, that allow for standard passage retrieval systems to work on content-only versions of the collection. That is, the offset and length are calculated over the text of the article, ignoring all mark-up. The evaluation measures are based directly on the highlighted passages, or arbitrary best-entry points, as identified by the assessors. As a result it is now possible to fairly compare systems retrieving elements, ranges of elements, or arbitrary passages. These changes address earlier requests to liberalize the retrieval format to ranges of elements [2] and later requests to liberalize to arbitrary passages of text [11].

The second main question is to compare focused retrieval directly to traditional article retrieval. Throughout the history of INEX, participating groups have found that article retrieval—a system retrieving the whole article by default—resulted in fairly competitive performance [e.g., 7, 10]. Note that every focused retrieval system also generates an underlying article ranking, simply by the order in which results from different articles are ranked. This is most clear in the Relevant in Context and Best in Context tasks, where the article ranking is an explicit part of the task description. To study the importance of the underlying article ranking quality, we derived article level judgments by treating every article with some highlighted text as relevant, derived article rankings from every submission on a first-come, first-served basis, and evaluated with standard measures. This will also shed light on the value of element or passage level evidence for document retrieval [1]. In addition to this, we also include queries derived from a proxy log in the topic set, and can derive judgments from the later clicks in the same proxy log, treating all clicked articles as relevant for the query at

hand. All submissions are also evaluated against these clicked Wikipedia pages, giving some insight in the differences between an IR test collection and real-world searching of Wikipedia.

The INEX 2008 Ad Hoc Track featured three tasks:

1. For the *Focused Task* a ranked-list of non-overlapping results (elements or passages) must be returned. It is evaluated at early precision relative to the highlighted (or believed relevant) text retrieved.
2. For the *Relevant in Context Task* non-overlapping results (elements or passages) must be returned, these are grouped by document. It is evaluated by mean average generalized precision where the generalized score per article is based on the retrieved highlighted text.
3. For the *Best in Context Task* a single starting point (element’s starting tag or passage offset) per article must be returned. It is also evaluated by mean average generalized precision but with the generalized score (per article) based on the distance to the assessor’s best-entry point.

We discuss the results for the three tasks, giving results for the top 10 participating groups and discussing the best scoring approaches in detail. We also examine the relative effectiveness of element and passage runs, and with content only (CO) queries and content and structure (CAS) queries.

The rest of the paper is organized as follows. First, Section 2 describes the INEX 2008 ad hoc retrieval tasks and measures. Section 3 details the collection, topics, and assessments of the INEX 2008 Ad Hoc Track. In Section 4, we report the results for the Focused Task (Section 4.2); the Relevant in Context Task (Section 4.3); and the Best in Context Task (Section 4.4). Section 5 details particular types of runs (such as CO versus CAS, and element versus passage), and on particular subsets of the topics (such as topics with a non-trivial CAS query). Section 6 looks at the article retrieval aspects of the submissions, both in terms of the judged topics treating any article with highlighted text as relevant, and in terms of clicked Wikipedia pages for queries derived from a proxy log. Finally, in Section 7, we discuss our findings and draw some conclusions.

2 Ad Hoc Retrieval Track

In this section, we briefly summarize the ad hoc retrieval tasks and the submission format (especially how elements and passages are identified). We also summarize the measures used for evaluation.

2.1 Tasks

Focused Task The scenario underlying the Focused Task is the return, to the user, of a ranked list of elements or passages for their topic of request. The Focused Task requires systems to find the most focused results that satisfy an information need, without returning “overlapping” elements (shorter is preferred in the case of equally relevant elements). Since ancestors elements and longer

passages are always relevant (to a greater or lesser extent) it is a challenge to choose the correct granularity.

The task has a number of assumptions:

Display the results are presented to the user as a ranked-list of results.

Users view the results top-down, one-by-one.

Relevant in Context Task The scenario underlying the Relevant in Context Task is the return of a ranked list of articles and within those articles the relevant information (captured by a set of non-overlapping elements or passages). A relevant article will likely contain relevant information that could be spread across different elements. The task requires systems to find a set of results that corresponds well to all relevant information in each relevant article. The task has a number of assumptions:

Display results will be grouped per article, in their original document order, access will be provided through further navigational means, such as a document heat-map or table of contents.

Users consider the article to be the most natural retrieval unit, and prefer an overview of relevance within this context.

Best in Context Task The scenario underlying the Best in Context Task is the return of a ranked list of articles and the identification of a best-entry-point from which a user should start reading each article in order to satisfy the information need. Even an article completely devoted to the topic of request will only have one best starting point from which to read (even if that is the beginning of the article). The task has a number of assumptions:

Display a single result per article.

Users consider articles to be natural unit of retrieval, but prefer to be guided to the best point from which to start reading the most relevant content.

2.2 Submission Format

Since XML retrieval approaches may return arbitrary results from within documents, a way to identify these nodes is needed. At INEX 2008, we allowed the submission of three types of results: XML elements; ranges of XML elements; and file-offset-length (FOL) text passages.

Element Results XML element results are identified by means of a file name and an element (node) path specification. File names in the Wikipedia collection are unique so that (with the .xml extension removed). The next example identifies 9996.xml as the target document from the Wikipedia collection.

```
<file>9996</file>
```

Element paths are given in XPath, but only fully specified paths are allowed. The next example identifies the first “article” element, then within that, the first “body” element, then the first “section” element, and finally within that the first “p” element.

```
<path>/article[1]/body[1]/section[1]/p[1]</path>
```

Importantly, XPath counts elements from 1 and counts element types. For example if a section had a title and two paragraphs then their paths would be: `title[1]`, `p[1]` and `p[2]`.

A result element, then, is identified unambiguously using the combination of file name and element path, as shown in the next example.

```
<result>
  <file>9996</file>
  <path>/article[1]/body[1]/section[1]/p[1]</path>
  <rsv>0.9999</rsv>
</result>
```

Ranges of Elements To support ranges of elements, elemental passages are given in the same format.¹ As a passage need not start and end in the same element, each is given separately. The following example is equivalent to the element result example above since it starts and ends on an element boundary.

```
<result>
  <file>9996</file>
  <passage start="/article[1]/body[1]/section[1]/p[1]"
    end="/article[1]/body[1]/section[1]/p[1]"/>
  <rsv>0.9999</rsv>
</result>
```

Note that this format is very convenient for specifying ranges of elements, e.g., the following example retrieves the first three sections.

```
<result>
  <file>9996</file>
  <passage start="/article[1]/body[1]/section[1]"
    end="/article[1]/body[1]/section[3]"/>
  <rsv>0.9999</rsv>
</result>
```

FOL passages Passage results can be given in File-Offset-Length (FOL) format, where offset and length are calculated in characters with respect to the textual content (ignoring all tags) of the XML file. A special text-only version of

¹ At INEX 2007, and in earlier qrels, an extended format allowing for optional character-offsets was used that allowed these passages to start or end in the middle of element or text-nodes. This format is superseded with the clean file-offset-length (FOL) passage format.

the collection is provided to facilitate the use of passage retrieval systems. File offsets start counting a 0 (zero).

The following example is effectively equivalent to the example element result above.

```
<result>
  <file>9996</file>
  <fol offset="461" length="202"/>
  <rsv>0.9999</rsv>
</result>
```

The paragraph starts at the 462th character (so 461 characters beyond the first character), and has a length of 202 characters.

2.3 Evaluation Measures

We briefly summarize the main measures used for the Ad Hoc Track. Since INEX 2007, we allow the retrieval of arbitrary passages of text matching the judges ability to regard any passage of text as relevant. Unfortunately this simple change has necessitated the deprecation of element-based metrics used in prior INEX campaigns because the “natural” retrieval unit is no longer an element, so elements cannot be used as the basis of measure. We note that properly evaluating the effectiveness in XML-IR remains an ongoing research question at INEX.

The INEX 2008 measures are solely based on the retrieval of highlighted text. We simplify all INEX tasks to highlighted text retrieval and assume that systems return all, and only, highlighted text. We then compare the characters of text retrieved by a search engine to the number and location of characters of text identified as relevant by the assessor. For best in context we use the distance between the best entry point in the run to that identified by an assessor.

Focused Task Recall is measured as the fraction of all highlighted text that has been retrieved. Precision is measured as the fraction of retrieved text that was highlighted. The notion of rank is relatively fluid for passages so we use an interpolated precision measure which calculates interpolated precision scores at selected recall levels. Since we are most interested in what happens in the first retrieved results, the INEX 2008 official measure is interpolated precision at 1% recall (iP[0.01]). We also present interpolated precision at other early recall points, and (mean average) interpolated precision over 101 standard recall points (0.00, 0.01, 0.02, ..., 1.00) as an overall measure.

Relevant in Context Task The evaluation of the Relevant in Context Task is based on the measures of generalized precision and recall [9], where the per document score reflects how well the retrieved text matches the relevant text in the document. Specifically, the per document score is the harmonic mean of precision and recall in terms of the fractions of retrieved and highlighted text

in the document. We use an F_β score with $\beta = 1/4$ making precision four times as important as recall (at INEX 2007, F_1 was used). We are most interested in overall performances so the main measure is mean average generalized precision (MAgP). We also present the generalized precision scores at early ranks (5, 10, 25, 50).

Best in Context Task The evaluation of the Best in Context Task is based on the measures of generalized precision and recall where the per document score reflects how well the retrieved entry point matches the best entry point in the document. Specifically, the per document score is a linear discounting function of the distance d (measured in characters)

$$\frac{n - d(x, b)}{n}$$

for $d < n$ and 0 otherwise. We use $n = 500$ which is roughly the number of characters corresponding to the visible part of the document on a screen (at INEX 2007, $n = 1,000$ was used). We are most interested in overall performance, and the main measure is mean average generalized precision (MAgP). We also show the generalized precision scores at early ranks (5, 10, 25, 50).

3 Ad Hoc Test Collection

In this section, we discuss the corpus, topics, and relevance assessments used in the Ad Hoc Track.

3.1 Corpus

The document collection was the Wikipedia XML Corpus based on the English Wikipedia in early 2006 [3]. The Wikipedia collection contains 659,338 Wikipedia articles. On average an article contains 161 XML nodes, where the average depth of a node in the XML tree of the document is 6.72.

The original Wiki syntax has been converted into XML, using both general tags of the layout structure (like *article*, *section*, *paragraph*, *title*, *list* and *item*), typographical tags (like *bold*, *emphatic*), and frequently occurring link-tags. For details see Denoyer and Gallinari [3].

3.2 Topics

The ad hoc topics were created by participants following precise instructions. Candidate topics contained a short CO (keyword) query, an optional structured CAS query, a one line description of the search request, and narrative with a details of the topic of request and the task context in which the information need arose. Figure 1 presents an example of an ad hoc topic. Based on the submitted candidate topics, 135 topics were selected for use in the INEX 2008 Ad Hoc Track as topic numbers 544–678.


```

<topic id="544" ct_no="6">
  <title>meaning of life</title>
  <castitle>
    //article[about(., philosophy)]//section[about(., meaning of life)]
  </castitle>
  <description>What is the meaning of life?</description>
  <narrative>
    I got bored of my life and started wondering what the meaning of
    life is. An element is relevant if it discusses the meaning of life
    from different perspectives, as long as it is serious. For example,
    Socrates discussing meaning of life is relevant, but something like
    "42" from H2G2 or "the meaning of life is cheese" from a comedy is
    irrelevant. An element must be self contained. An element that is a
    list of links is considered irrelevant because it is not
    self-contained in the sense that I don't know in which context the
    links are given.
  </narrative>
</topic>

```

Fig. 1. INEX 2008 Ad Hoc Track topic 544.

In addition, 150 queries were derived from a proxy-log for use in the INEX 2008 Ad Hoc Track as topic numbers 679–828. For these topics, as well as the candidate topics without a `<castitle>` field, a default CAS-query was added based on the CO-query: `/**[about(., "CO-query")]`.

3.3 Judgments

Topics were assessed by participants following precise instructions. The assessors used the new GPXrai assessment system that assists assessors in highlight relevant text. Topic assessors were asked to mark all, and only, relevant text in a pool of documents. After assessing an article with relevance, a separate best entry point decision was made by the assessor. The Focused and Relevant in Context Tasks were evaluated against the text highlighted by the assessors, whereas the Best in Context Task was evaluated against the best-entry-points.

The relevance judgments were frozen on October 22, 2008. At this time 70 topics had been fully assessed. Moreover, 11 topics were judged by two separate assessors, each without the knowledge of the other. All results in this paper refer to the 70 topics with the judgments of the first assigned assessor, which is typically the topic author.

- The 70 assessed topics were: 544–547, 550–553, 555–557, 559, 561, 562–563, 565, 570, 574, 576–582, 585–587, 592, 595–598, 600–603, 607, 609–611, 613, 616–617, 624, 626, 628, 629, 634–637, 641–644, 646–647, 649–650, 656–657, 659, 666–669, 673, 675, and 677.

In addition, there are clicked Wikipedia pages available in the proxy log for 125 topics:

Table 1. Statistics over judged and relevant articles per topic.

	total		# per topic				
	topics	number	min	max	median	mean	st.dev
judged articles	70	42,105	502	618	603	601.5	15.3
articles with relevance	70	4,850	2	375	49	69.3	68.7
highlighted passages	70	7,510	3	906	59	107.3	131.0
highlighted characters	70	11,337,505	1,419	1,113,578	99,569	161,964.4	132,544.9
Unique articles with clicks	125	225	1	10	1	1.8	1.5
Total clicked articles	125	532	1	24	3	4.3	3.8

Fig. 2. Distribution of passages over articles.

- The 125 topics with clicked articles are numbered: 679–682, 684–685, 687–693, 695–704, 706–708, 711–727, 729–732, 734–751, 753–776, 778, 780–782, 784, 786–787, 789–790, 792–793, 795–796, 799–804, 806–807, 809–810, 812–813, 816–819, 821–824, and 826–828.

Table 1 presents statistics of the number of judged and relevant articles, and passages. In total 42,105 articles were judged. Relevant passages were found in 4,850 articles. The mean number of relevant articles per topic is 69, but the distribution is skewed with a median of 49. There were 7,510 highlighted passages. The mean was 107 passages and the median was 59 passages per topic.²

Table 1 also includes some statistics of the number of clicked articles in the proxy log. There are in total 225 clicked articles (unique per topic) over in total 125 topics, with a mean of 1.8 and a median of 1 clicked article per topic. We filtered the log for queries issued by multiple persons, and can also count the total number of clicks. Here, we see a total of 532 clicks (on the same 225 articles before), with a mean of 4.3 and a median of 3 clicks per topic. It is clear that the topics and clicked articles from the log are very different in character from the ad hoc topics.

² Recall from above that for the Focused Task the main effectiveness measures is precision at 1% recall. Given that the average topic has 107 relevant passages in 69 articles, the 1% recall roughly corresponds to a relevant passage retrieved—for many systems this will be accomplished by the first or first few results.

Table 2. Statistics over best entry point judgement.

	#	topics number	min	max	median	mean	st.dev
best entry point offset	70	4,850	1	87,982	14	1,746.2	4,826.5
first relevant character offset	70	4,850	1	87,982	20	1,821.4	4,862.9
fraction highlighted text	70	4,850	0.0005	1.000	0.580	0.549	0.425

Fig. 3. Distribution of best entry point offsets.

Figure 2 presents the number of articles with the given number of passages. The vast majority of relevant articles (3,696 out of 4,850) had only a single highlighted passage, and the number of passages quickly tapers off.

Assessors were requested to provide a separate best entry point (BEP) judgement, for every article where they highlighted relevant text. Table 2 presents statistics on the best entry point offset, on the first highlighted or relevant character, and on the fraction of highlighted text in relevant articles. We first look at the BEPs. The mean BEP is well within the article with 1,746 but the distribution is very skewed with a median BEP offset of only 14. Figure 3 shows the distribution of the character offsets of the 4,850 best entry points. It is clear that the overwhelming majority of BEPs is at the beginning of the article.

The statistics of the first highlighted or relevant character (FRC) in Table 2 give very similar numbers as the BEP offsets: the mean offset of the first relevant character is 1,821 but the median offset is only 20. This suggests a relation between the BEP offset and the FRC offset. Figure 4 shows a scatter plot the BEP and FRC offsets. Two observations present themselves. First, there is a clear diagonal where the BEP is positioned exactly at the first highlighted character in the article. Second, there is also a vertical line at BEP offset zero, indicating a tendency to put the BEP at the start of the article even when the relevant text appears later on.

Finally, the statistics on the fraction of highlighted text in Table 2 show that amount of relevant text varies from almost nothing to almost everything. The mean fraction is 0.55, and the median is 0.58, indicating that typically over half the article is relevant. Given that the majority of relevant articles contain such a large fraction of relevant text plausibly explains that BEPs being frequently positioned on or near the start of the article.

Fig. 4. Scatter plot of best entry point offsets versus the first relevant character.

Table 3. Candidate Topic Questionnaire.

B1	How familiar are you with the subject matter of the topic?
B2	Would you search for this topic in real-life?
B3	Does your query differ from what you would type in a web search engine?
B4	Are you looking for very specific information?
B5	Are you interested in reading a lot of relevant information on the topic?
B6	Could the topic be satisfied by combining the information in different (parts of) documents?
B7	Is the topic based on a seen relevant (part of a) document?
B8	Can information of equal relevance to the topic be found in several documents?
B9	Approximately how many articles in the whole collection do you expect to contain relevant information?
B10	Approximately how many relevant document parts do you expect in the whole collection?
B11	Could a relevant result be (check all that apply): a single sentence; a single paragraph; a single (sub)section; a whole article
B12	Can the topic be completely satisfied by a single relevant result?
B13	Is there additional value in reading several relevant results?
B14	Is there additional value in knowing all relevant results?
B15	Would you prefer seeing: only the best results; all relevant results; don't know
B16	Would you prefer seeing: isolated document parts; the article's context; don't know
B17	Do you assume perfect knowledge of the DTD?
B18	Do you assume that the structure of at least one relevant result is known?
B19	Do you assume that references to the document structure are vague and imprecise?
B20	Comments or suggestions on any of the above (optional)

3.4 Questionnaires

At INEX 2008, all candidate topic authors and assessors were asked to complete a questionnaire designed to capture the context of the topic author and the topic of request. The candidate topic questionnaire (shown in Table 3) featured 20 questions capturing contextual data on the search request. The post-assessment

Table 4. Post Assessment Questionnaire.

-
- C1 Did you submit this topic to INEX?
 - C2 How familiar were you with the subject matter of the topic?
 - C3 How hard was it to decide whether information was relevant?
 - C4 Is Wikipedia an obvious source to look for information on the topic?
 - C5 Can a highlighted passage be (check all that apply): a single sentence; a single paragraph; a single (sub)section; a whole article
 - C6 Is a single highlighted passage enough to answer the topic?
 - C7 Are highlighted passages still informative when presented out of context?
 - C8 How often does relevant information occur in an article about something else?
 - C9 How well does the total length of highlighted text correspond to the usefulness of an article?
 - C10 Which of the following two strategies is closer to your actual highlighting:
 - (I) I located useful articles and highlighted the best passages and nothing more,
 - (II) I highlighted all text relevant according to narrative, even if this meant highlighting an entire article.
 - C11 Can a best entry point be (check all that apply): the start of a highlighted passage; the sectioning structure containing the highlighted text; the start of the article
 - C12 Does the best entry point correspond to the best passage?
 - C13 Does the best entry point correspond to the first passage?
 - C14 Comments or suggestions on any of the above (optional)

questionnaire (shown in Table 4) featured 14 questions capturing further contextual data on the search request, and the way the topic has been judged (a few questions on GPXrai were added to the end).

The responses to the questionnaires show a considerable variation over topics and topic authors in terms of topic familiarity; the type of information requested; the expected results; the interpretation of structural information in the search request; the meaning of a highlighted passage; and the meaning of best entry points. There is a need for further analysis of the contextual data of the topics in relation to the results of the INEX 2008 Ad Hoc Track.

4 Ad Hoc Retrieval Results

In this section, we discuss, for the three ad hoc tasks, the participants and their results.

4.1 Participation

A total of 163 runs were submitted by 23 participating groups. Table 5 lists the participants and the number of runs they submitted, also broken down over the tasks (Focused, Relevant in Context, or Best in Context); the used query (Content-Only or Content-And-Structure); and the used result type (Element, Passage or FOL). Unfortunately, no less than 27 runs turned out to be invalid and will only be evaluated with respect to their “article retrieval” value in Section 6.

Participants were allowed to submit up to three element result-type runs per task and three passage result-type runs per task (for all three tasks). This

Table 5. Participants in the Ad Hoc Track.

Id Participant	Focused	Relevant in Context	Best in Context	CO query	CAS query	Element results	Passage results	FOL results	# valid runs	# submitted runs
4 University of Otago	0	6	0	6	0	3	3	0	6	6
5 Queensland University of Technology	6	6	6	15	3	9	9	0	18	18
6 University of Amsterdam	6	6	3	9	6	13	0	2	15	15
9 University of Helsinki	3	0	0	3	0	3	0	0	3	3
10 Max-Planck-Institut Informatik	3	1	1	5	0	5	0	0	5	5
12 University of Granada	3	3	3	9	0	9	0	0	9	9
14 University of California, Berkeley	2	0	1	3	0	3	0	0	3	3
16 University of Frankfurt	1	3	3	0	7	7	0	0	7	9
22 ENSM-SE	2	0	0	2	0	0	0	2	2	9
25 Renmin University of China	3	0	1	2	2	4	0	0	4	4
29 INDIAN STATISTICAL INSTITUTE	3	0	0	3	0	3	0	0	3	3
37 Katholieke Universiteit Leuven	6	0	0	3	3	6	0	0	6	6
40 IRIT	0	0	2	1	1	2	0	0	2	6
42 University of Toronto	2	0	0	0	2	2	0	0	2	3
48 LIG	3	2	0	5	0	5	0	0	5	5
55 Doshisha University	0	0	1	0	1	1	0	0	1	3
56 JustSystems Corporation	3	3	3	6	3	9	0	0	9	9
60 Saint Etienne University	3	0	0	3	0	3	0	0	3	9
61 Universit Libre de Bruxelles	0	0	0	0	0	0	0	0	0	2
68 University Pierre et Marie Curie - LIP6	2	0	0	2	0	2	0	0	2	2
72 University of Minnesota Duluth	2	2	2	6	0	6	0	0	6	6
78 University of Waterloo	3	3	4	10	0	8	2	0	10	13
92 University of Lyon3	5	5	5	15	0	15	0	0	15	15
Total runs	61	40	35	108	28	118	14	4	136	163

totaled to 18 runs per participant.³ The submissions are spread well over the ad hoc retrieval tasks with 61 submissions for Focused, 40 submissions for Relevant in Context, and 35 submissions for Best in Context.

4.2 Focused Task

We now discuss the results of the Focused Task in which a ranked-list of non-overlapping results (elements or passages) was required. The official measure for the task was (mean) interpolated precision at 1% recall (iP[0.01]). Table 6 shows the best run of the top 10 participating groups. The first column gives the

³ As it turns out, two groups submitted more runs than allowed: *University of Lyon3* submitted 6 extra element runs, and *University of Amsterdam* submitted 4 extra element runs. At this moment, we have not decided on any repercussions other than mentioning them in this footnote.

Table 6. Top 10 Participants in the Ad Hoc Track Focused Task.

Participant	iP[.00]	iP[.01]	iP[.05]	iP[.10]	MAiP
p78-FOERStep	0.7657	0.6873	0.5700	0.4879	0.2071
p10-TOPXCOarti	0.6804	0.6795	0.5807	0.5265	0.2967
p48-LIGMLFOCRI	0.7114	0.6665	0.5210	0.4216	0.1441
p92-manualQEin*	0.6664	0.6664	0.6139	0.5540	0.3065
p60-JMUexpe142	0.6918	0.6640	0.5800	0.4986	0.2342
p9-UHelRun394	0.7109	0.6619	0.5532	0.5028	0.2251
p14-T2FBCOPARA	0.7319	0.6395	0.4906	0.4026	0.1392
p25-weightedfi	0.6553	0.6346	0.5490	0.5222	0.2647
p5-GPX1COFOCe	0.6818	0.6344	0.5693	0.5178	0.2587
p29-LMnofb020	0.6830	0.6337	0.5537	0.5100	0.2847

participant, see Table 5 for the full name of group. The second to fifth column give the interpolated precision at 0%, 1%, 5%, and 10% recall. The sixth column gives mean average interpolated precision over 101 standard recall levels (0%, 1%, ..., 100%).

Here we briefly summarize what is currently known about the experiments conducted by the top five groups (based on official measure for the task, iP[0.01]).

University of Waterloo Element retrieval run using the CO query. Description: the run uses the Okapi BM25 model in Wumpus to score all content-bearing elements such as sections and paragraphs using Okapi BM25. In addition, scores were boosted by doubling the tf values of the first 10 words of an element.

Max-Planck-Institut für Informatik Element retrieval run using the CO query. Description: The TopX system retrieving only article elements, using a linear combination of a BM25 content score with a BM25 proximity score that also takes document structure into account.

LIG Grenoble An element retrieval run using the CO query. Description: Based on a language Model using a Dirichlet smoothing, and equally weighting element score and its context score, where the context score are based on the collection-links in Wikipedia.

University of Lyon3 A *manual* element retrieval run using the CO query. Description: Using indri search engine in Lemur with manually expanded queries from CO, description and narrative fields. The run is retrieving only articles.

Saint Etienne University An element retrieval run using the CO query. Description: A probabilistic model used to evaluate a weight for each tag: "the probability that tags distinguishes terms which are the most relevant", i.e. based on the fact that the tag contains relevant or non relevant passages. The resulting tag weights are incorporated into an element-level run with BM25 weighting.

Based on the information from these and other participants:

- All ten runs use the CO query. The fourth run, *p92-manualQEin*, uses a manually expanded query using words from the description and narrative

Table 7. Top 10 Participants in the Ad Hoc Track Relevant in Context Task.

Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAgP
p78-RICBest	0.4052	0.3414	0.2739	0.2184	0.2263
p5-GPX1CORICe	0.3737	0.3430	0.2658	0.2135	0.2097
p92-manualQEin*	0.4138	0.3564	0.2659	0.2069	0.2092
p10-TOPXCOallA	0.3650	0.3049	0.2352	0.1908	0.1933
p4-WHOLEDOC	0.3696	0.3242	0.2476	0.1944	0.1917
p6-inex08artB	0.3481	0.2991	0.2200	0.1726	0.1752
p72-UMDRic2	0.3828	0.3341	0.2342	0.1882	0.1714
p12-p8u3exp511	0.2933	0.2710	0.2154	0.1612	0.1575
p56-VSMRIP05	0.3281	0.2638	0.2097	0.1608	0.1495
p48-LIGMLRIC4O	0.3595	0.3069	0.2303	0.1708	0.1486

fields. The tenth run, *p29-LMnofb020*, is an automatic run using the title and description fields. All other runs use only the CO query in the title field.

- All runs retrieve elements as results.
- The systems at rank second (*p10-TOPXCOarti*), fourth (*p92-manualQEin*), and tenth (*p29-LMnofb020*), are retrieving only full articles.

4.3 Relevant in Context Task

We now discuss the results of the Relevant in Context Task in which non-overlapping results (elements or passages) need to be returned grouped by the article they came from. The task was evaluated using generalized precision where the generalized score per article was based on the retrieved highlighted text. The official measure for the task was mean average generalized precision (MAgP).

Table 7 shows the top 10 participating groups (only the best run per group is shown) in the Relevant in Context Task. The first column lists the participant, see Table 5 for the full name of group. The second to fifth column list generalized precision at 5, 10, 25, 50 retrieved articles. The sixth column lists mean average generalized precision.

Here we briefly summarize the information available about the experiments conducted by the top five groups (based on MAgP).

University of Waterloo Element retrieval run using the CO query. Description: the run uses the Okapi BM25 model in Wumpus to score all content-bearing elements such as sections and paragraphs using Okapi BM25, and grouped the results by articles and ranked the articles by their best scoring element.

Queensland University of Technology Element retrieval run using the CO query. Description: GPX run using a `/*[about(. , keywords)]` query, serving non-overlapping elements grouped per article, with the articles ordered by their best scoring element.

University of Lyon3 A manual element retrieval run using the CO query. Description: the same as the Focused run above. In fact it is literally the same article ranking as the Focused run. Recall that the run is retrieving only whole articles.

Table 8. Top 10 Participants in the Ad Hoc Track Best in Context Task.

Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAgP
p78-BICER	0.3811	0.3233	0.2498	0.1979	0.2207
p92-manualQEin*	0.4087	0.3645	0.2784	0.2203	0.2171
p25-weightedfi	0.3454	0.3003	0.2481	0.2004	0.2009
p5-GPX1COBICp	0.3655	0.3367	0.2572	0.2019	0.1970
p6-submitinex	0.3447	0.2870	0.2203	0.1681	0.1693
p10-TOPXCOallB	0.2392	0.2321	0.1875	0.1528	0.1689
p12-p8u3exp501	0.2490	0.2303	0.1935	0.1487	0.1457
p72-UMDBIC1	0.3163	0.2710	0.1857	0.1451	0.1438
p56-VSMRIP08	0.2269	0.2010	0.1736	0.1397	0.1311
p40-xfirmcos07	0.2402	0.1869	0.1347	0.1083	0.0951

Max-Planck-Institut für Informatik Element retrieval run using the CO query. Description: An element retrieval run using the new BM25 scoring function (i.e., considering each element as “document” and then computing a standard BM25 model), selecting non-overlapping elements based on score, and grouping them per article with the articles ranked by their highest scoring element.

University of Otago Element retrieval run using the CO query. Description: BM25 is used to select and rank the top 1,500 documents and whole documents are selected as the passage. That is, the run is retrieving only whole articles.

Based on the information from these and other participants:

- The runs ranked sixth (*p6-inex08artB*) and ninth (*p56-VSMRIP05*) are using the CAS query. The run ranked third, *p92-manualQEin*, is using a manually expanded query based on keywords in the description and narrative. All other runs use only the CO query in the topic’s title field.
- All runs retrieve elements as results.
- Solid article ranking seems a prerequisite for good overall performance, with third best run, *p92-manualQEin*, the fifth best run, *p4-WHOLEDOC*, and the ninth best run, *p56-VSMRIP05*, retrieving only full articles.

4.4 Best in Context Task

We now discuss the results of the Best in Context Task in which documents were ranked on topical relevance and a single best entry point into the document was identified. The Best in Context Task was evaluated using generalized precision but here the generalized score per article was based on the distance to the assessor’s best-entry point. The official measure for the task was mean average generalized precision (MAgP).

Table 8 shows the top 10 participating groups (only the best run per group is shown) in the Best in Context Task. The first column lists the participant, see Table 5 for the full name of group. The second to fifth column list generalized

precision at 5, 10, 25, 50 retrieved articles. The sixth column lists mean average generalized precision.

Here we briefly summarize the information available about the experiments conducted by the top five groups (based on MAGP).

University of Waterloo Element retrieval run using the CO query. Description: the run uses the Okapi BM25 model in Wumpus to score all content-bearing elements such as sections and paragraphs using Okapi BM25, and kept only the best scoring element per article.

University of Lyon3 A manual element retrieval run using the CO query. Description: the same as the Focused and Relevant in Context runs above. In fact all three runs have literally the same article ranking. This run is retrieving the start of the whole article as best entry point, in other words an article retrieval run.

Renmin University of China Element retrieval run using the CO query. Description: using language model to compute RSV at leaf level combined with aggregation at retrieval time, assuming independence.

Queensland University of Technology Run retrieving ranges of elements using the CO query. The run is always returning a whole article, setting the BEP at the very start of the article. Description: GPX run using a `//*[about(.,keywords)]` query, ranking articles by their best scoring element, but transformed to return the complete article as a passages. This is effectively an article level GPX run.

University of Amsterdam Run retrieving FOL passages using the CO query. Description: language model with local indegree prior, setting the BEP always at the start of the article. Since the offset is always zero, this is similar to an article retrieval run.

Based on the information from these and other participants:

- As for the Relevant in Context Task, we see again that solid article ranking is very important. In fact, we see runs putting the BEP at the start of all the retrieved articles at rank two (*p92-manualQEIn*), rank four (*p5-GPX1COBICp*), and rank five (*p6-submitinex*).
- The fourth ranked run, *p5-GPX1COBICp*, uses ranges of elements, albeit a degenerate case where always the full article is selected. The fifth run, *p6-submitinex*, uses fol passages, albeit again a degenerate case where the BEP is always the zero offset.
- With the exception of the runs ranked nine (*p56-VSMRIP08*) and ten (*p40-xfirmcos07*), which used the CAS query, all the other best runs per group use the CO query.

4.5 Significance Tests

We tested whether higher ranked systems were significantly better than lower ranked system, using a t-test (one-tailed) at 95%. Table 9 shows, for each task, whether it is significantly better (indicated by “★”) than lower ranked runs. For

Table 9. Statistical significance (t-test, one-tailed, 95%).

(a) Focused Task											(b) Relevant in Context Task											(c) Best in Context Task										
	1	2	3	4	5	6	7	8	9	10		1	2	3	4	5	6	7	8	9	10		1	2	3	4	5	6	7	8	9	10
p78	-	-	-	-	-	-	-	-	-	-	p78	*	-	*	*	*	*	*	*	*	*	p78	-	-	*	*	*	*	*	*	*	*
p10	-	-	-	-	-	-	-	-	-	-	p5	-	*	-	*	*	*	*	*	*	p92	-	-	*	*	*	*	*	*	*	*	*
p48	-	-	-	-	-	-	-	-	-	-	p92	-	-	*	*	*	*	*	*	*	p25	-	*	*	*	*	*	*	*	*	*	*
p92	-	-	-	-	-	-	-	-	-	-	p10	-	-	*	*	*	*	*	*	*	p5	*	-	*	*	*	*	*	*	*	*	*
p60	-	-	-	-	-	-	-	-	-	-	p4	-	-	*	*	*	*	*	*	*	p6	-	-	*	*	*	*	*	*	*	*	*
p9	-	-	-	-	-	-	-	-	-	-	p6	-	-	*	*	*	*	*	*	*	p10	-	-	*	*	*	*	*	*	*	*	*
p14	-	-	-	-	-	-	-	-	-	-	p72	-	-	*	*	*	*	*	*	*	p12	-	-	*	*	*	*	*	*	*	*	*
p25	-	-	-	-	-	-	-	-	-	-	p12	-	-	*	*	*	*	*	*	*	p72	-	-	*	*	*	*	*	*	*	*	*
p5	-	-	-	-	-	-	-	-	-	-	p56	-	-	*	*	*	*	*	*	*	p56	-	-	*	*	*	*	*	*	*	*	*
p29	-	-	-	-	-	-	-	-	-	-	p48	-	-	*	*	*	*	*	*	*	p40	-	-	*	*	*	*	*	*	*	*	*

example, For the Focused Task, we see that the early precision (at 1% recall) is a rather unstable measure and none of the runs are significantly different. Hence we should be careful when drawing conclusions based on the Focused Task results. For the Relevant in Context Task, we see that the top run is significantly better than ranks 2 and 4 through 10, the second best run better than ranks 4 and 6 through 10, the third ranked system better than ranks 6 through 10, and the fourth and fifth ranked systems better than ranks 8 through 10. For the Best in Context Task, we see that the top run is significantly better than ranks 4 through 10, the second and third runs significantly better than ranks 5 to 10. The fourth ranked system is better than the systems ranked 5 and 7 to 10, and the fifth ranked system better than ranks 9 and 10.

5 Analysis of Run and Topic Types

In this section, we will discuss relative effectiveness of element and passage retrieval approaches, and on the relative effectiveness of systems using the keyword and structured queries.

5.1 Elements versus passages

We received 18 submissions using ranges of elements of FOL-passage results, from in total 5 participating groups. We will look at the relative effectiveness of element and passage runs.

As we saw above, in Section 4, for all three tasks the best scoring runs used elements as the unit of retrieval. Table 10 shows the best runs using ranges of elements or FOL passages for the three ad hoc tasks. All these runs use the CO query. As it turns out, the best focused run using passages ranks outside the top scoring runs in Table 6; the best relevant in context run using passages is ranked fifth among the top scoring runs in Table 7; and the best best in context run using passages is ranked fourth among the top scoring runs in Table 8. This outcome is consistent with earlier results using passage-based element retrieval,

Table 10. Ad Hoc Track: Runs with ranges of elements or FOL passages.

(a) Focused Task					
Participant	iP[.00]	iP[.01]	iP[.05]	iP[.10]	MAiP
p5-GPX2COFOCP	0.6308	0.6301	0.5379	0.4699	0.2502
p22-EMSEFocus*	0.6745	0.5703	0.4481	0.3845	0.1545

(b) Relevant in Context Task					
Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAGP
p4-WHOLEDOCPA	0.3696	0.3242	0.2476	0.1944	0.1917
p5-GPX1CORICp	0.3544	0.3209	0.2413	0.1860	0.1891

(c) Best in Context Task					
Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAGP
p5-GPX1COBICp	0.3655	0.3367	0.2572	0.2019	0.1970
p6-submitinex	0.3447	0.2870	0.2203	0.1681	0.1693
p78-BICPRplus	0.2585	0.2206	0.1642	0.1248	0.1237

where passage retrieval approaches showed comparable but not superior behavior to element retrieval approaches [4, 5].

However, looking at the runs in more detail, their character is often unlike what one would expect from a “passage” retrieval run. For Focused, *p5-GPX2COFOCP* is an article retrieving run using ranges of elements; and *p22-EMSEFocus* is a manual query run using FOL passages. For Relevant in Context, both *p4-WHOLEDOCPA* and *p5-GPX1CORICp* are article retrieving runs using ranges of elements. For Best in Context, *p5-GPX1COBICp* is an article runs using ranges of elements; *p6-submitinex* is an article run using FOL passages; and *p78-BICPRplus* is an element retrieving run using ranges of elements. So, all but two of the runs retrieve only articles. Hence, this is not sufficient evidence to warrant any conclusion on the effectiveness of passage level results. We hope and expect that the test collection and the passage runs will be used for further research into the relative effectiveness of element and passage retrieval approaches.

5.2 CO versus CAS

We now look at the relative effectiveness of the keyword (CO) and structured (CAS) queries. As we saw above, in Section 4, one of the best runs per group for the Relevant in Context Task, and two of the top 10 runs for the Best in Context Task used the CAS query.

All topics have a CAS query since artificial CAS queries of the form

```
/*[about(., keyword title)]
```

were added to topics without CAS title. Table 11 show the distribution of target elements. In total 86 topics had a non-trivial CAS query.⁴ These CAS topics

⁴ Note that some of the wild-card topics (using the “*” target) in Table 11 had non-trivial about-predicates and hence have not been regarded as trivial CAS queries.

Table 11. CAS query target elements over all 135 topics.

Target Element	Frequency
*	51
section	39
article	30
p	11
figure	3
body	1

are numbered 544–550, 553–556, 564, 567, 568, 572, 574, 576–578, 580, 583, 584, 586–591, 597–605, 607, 608, 610, 615–625, 627, 629–633, 635–640, 646, 651–655, 658, 659, 661–670, 673, and 675–678. As it turned out, 39 of these CAS topics were assessed. The results presented here are restricted to only these 39 CAS topics.

Table 12 lists the top 10 participants measured using just the 39 CAS topics and for the Focused Task (a), the Relevant in Context Task (b), and the Best in Context Task (c). For the Focused Task the CAS runs score lower than the CO query runs. For the Relevant in Context Task, the best CAS run would have ranked fifth among the CO runs. For the Best in Context Task, the best CAS run would rank seventh among the CO runs. Overall, we see that teams submitting runs with both types of queries have higher scoring CO runs, with participant 6 as a notable exception for Relevant in Context.

6 Analysis of Article Retrieval

In this section, we will look in detail at the effectiveness of Ad Hoc Track submissions as article retrieval systems. We look first at the article rankings in terms of the Ad Hoc Track judgements—treating every article that contains highlighted text as relevant. Then, we look at the article rankings in terms of the clicked pages for the topics derived from the proxy log—treating every clicked articles as relevant.

6.1 Article retrieval: Relevance Judgments

We will first look at the topics judged during INEX 2008, the same topics as in earlier sections, but now using the judgments to derive standard document-level relevance by regarding an article as relevant if some part of it is highlighted by the assessor. Throughout this section, we derive an article retrieval run from every submission using a first-come, first served mapping. That is, we simply keep every first occurrence of an article (retrieved indirectly through some element contained in it) and ignore further results from the same article.

We use `trec_eval` to evaluate the mapped runs and qrels, and use mean average precision (map) as the main measure. Since all runs are now article retrieval runs, the differences between the tasks disappear. Moreover, runs violating the

Table 12. Ad Hoc Track CAS Topics: CO runs (left-hand side) versus CAS runs (right-hand side).

(a) Focused Task										
Participant	iP[.00]	iP[.01]	iP[.05]	iP[.10]	MAiP	Participant	iP[.00]	iP[.01]	iP[.05]	iP[.10] MAiP
p60-JMUexpe136	0.7321	0.7245	0.6416	0.5861	0.2926	p6-inex08artB	0.6514	0.6379	0.5947	0.5072 0.2255
p48-LIGMLFOCR1	0.7496	0.7209	0.5307	0.4440	0.1569	p56-VSMRIP02	0.7515	0.6314	0.4760	0.3656 0.1396
p78-FOER	0.7263	0.7064	0.6070	0.5470	0.2222	p5-GPX3COSFOC	0.6232	0.6220	0.5521	0.4617 0.2134
p5-GPX1COFOCe	0.7167	0.6970	0.6416	0.5607	0.2613	p25-RUCLLP08	0.5969	0.5969	0.5815	0.5439 0.2486
p29-LMnofb020	0.7193	0.6759	0.5919	0.5553	0.2938	p37-kulcaselem	0.6817	0.5611	0.3525	0.2720 0.1256
p10-TOPXCOallF	0.7482	0.6657	0.5514	0.4872	0.1923	p42-B2U0visith	0.6030	0.5360	0.4823	0.4442 0.1736
p25-weightedfi	0.6665	0.6634	0.5915	0.5588	0.2662	p16-001RunofUn	0.3110	0.2268	0.1673	0.1206 0.0364
p6-inex08artB	0.6689	0.6571	0.5397	0.4941	0.2098					
p9-UHelRun394	0.7024	0.6515	0.5555	0.5189	0.2246					
p72-UMDFocused	0.7206	0.6386	0.4913	0.3812	0.1111					

(b) Relevant in Context Task										
Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAgP	Participant	gP[5]	gP[10]	gP[25]	gP[50] MAgP
p78-RICBest	0.4760	0.3769	0.2975	0.2261	0.2476	p6-inex08artB	0.3799	0.3292	0.2379	0.1831 0.1932
p5-GPX1CORICe	0.3946	0.3518	0.2660	0.2156	0.2161	p5-GPX3COSRIC	0.3482	0.3232	0.2381	0.1918 0.1762
p4-WHOLEDOC	0.3936	0.3492	0.2491	0.1991	0.2114	p56-VSMRIP05	0.3401	0.2796	0.2133	0.1610 0.1498
p10-TOPXCOallA	0.3892	0.3170	0.2339	0.1897	0.1963	p16-009RunofUn	0.0153	0.0156	0.0123	0.0095 0.0023
p92-manualQEin*	0.3767	0.3370	0.2474	0.1943	0.1920					
p6-inex08artB	0.3711	0.3114	0.2274	0.1778	0.1892					
p72-UMDRic2	0.3908	0.3412	0.2280	0.1858	0.1740					
p12-p8u3exp511	0.3178	0.2855	0.2227	0.1622	0.1673					
p48-LIGMLRIC4O	0.3769	0.3384	0.2451	0.1827	0.1581					
p56-VSMRIP04	0.2264	0.2005	0.1664	0.1345	0.1263					

(c) Best in Context Task										
Participant	gP[5]	gP[10]	gP[25]	gP[50]	MAgP	Participant	gP[5]	gP[10]	gP[25]	gP[50] MAgP
p78-BICER	0.3883	0.3361	0.2514	0.1931	0.2162	p5-GPX3COSBIC	0.3109	0.2883	0.2235	0.1780 0.1659
p25-weightedfi	0.3342	0.3016	0.2360	0.1934	0.1992	p56-VSMRIP08	0.2123	0.1911	0.1481	0.1214 0.1227
p5-GPX1COBICp	0.3663	0.3358	0.2494	0.1911	0.1977	p40-xfirmcos07	0.2381	0.1794	0.1348	0.1078 0.0908
p92-manualQEin*	0.3677	0.3357	0.2558	0.2052	0.1939	p55-KikoriBest	0.1817	0.1721	0.1422	0.1123 0.0803
p10-TOPXCOallB	0.2424	0.2397	0.1769	0.1447	0.1723	p16-006RunofUn	0.0307	0.0347	0.0307	0.0261 0.0128
p6-submitinex	0.3454	0.3037	0.2248	0.1698	0.1707					
p12-p8u3exp501	0.2536	0.2373	0.1924	0.1412	0.1440					
p72-UMDBIC1	0.3171	0.2725	0.1746	0.1366	0.1363					
p56-VSMRIP09	0.1562	0.1537	0.1377	0.1122	0.1034					
p40-xfirmbicco	0.1594	0.1521	0.1357	0.1122	0.0656					

task requirements—most notably non-overlapping results for all tasks, and having scattered results from the same article in relevant in context—are now also considered, and we work with all 163 runs submitted to the Ad Hoc Track.

Table 13 shows the best run of the top 10 participating groups. The first column gives the participant, see Table 5 for the full name of group. The second

Table 13. Top 10 Participants in the Ad Hoc Track: Article retrieval.

Participant	P5	P10	1/rank	map	bpref
p78-BICER	0.6200	0.5257	0.8711	0.3753	0.3693
p92-manualQEin*	0.6371	0.5843	0.8322	0.3601	0.3917
p10-TOPXCOarti	0.5914	0.5386	0.8635	0.3489	0.3620
p5-GPX1COBICe	0.5686	0.5214	0.7868	0.3390	0.3580
p37-kulcoeleme	0.5229	0.4500	0.7468	0.3240	0.3335
p25-weightedfi	0.4914	0.4600	0.7192	0.3224	0.3350
p60-JMUexpe136	0.5429	0.4814	0.7843	0.3173	0.3380
p29-VSMfbElts0	0.5486	0.4800	0.7955	0.3163	0.3385
p9-UHelRun293	0.5714	0.4957	0.7766	0.3113	0.3317
p4-SWKL200	0.5657	0.4943	0.7950	0.3086	0.3292

and third column give the precision at ranks 5 and 10, respectively. The fourth column gives the mean reciprocal rank. The fifth column gives mean average precision. The sixth column gives binary preference measures (using the top R judged non-relevant documents). Recall from the above that second ranked run (*p92-manualQEin*) is a manual article retrieval run submitted to all three tasks. Also the run ranked three (*p10-TOPXCOarti*) and the run ranked seven (*p60-JMUexpe136*) retrieve exclusively articles. The relative effectiveness of these article retrieval runs in terms of their article ranking is no surprise. Furthermore, we see submissions from all three ad hoc tasks. Most notably runs from the Best in Context task at ranks 1, 2, 4, and 6; runs from the Focused task at ranks 2, 3, 5, 7, 8, and 9; and runs from the Relevant in Context task at ranks 2 and 10.

If we break-down all runs over the original tasks, shown on the left-hand side of Table 14, we can compare the ranking to Section 4 above. We see some runs that are familiar from the earlier tables: three Focused runs correspond to Table 6, five Relevant in Context runs correspond to Table 7, and seven Best in Context runs correspond to Table 8. More formally, we looked at how the two system rankings correlate using kendall’s tau.

- Over all 61 Focused task submissions the system rank correlation is 0.526 between $iP[0.01]$ and map, and 0.574 between MAiP and map.
- Over all 35 Relevant in Context submissions the system rank correlation between MAgP and map is 0.792.
- Over all 40 Best in Context submissions the system rank correlation is 0.787

Overall, we see a reasonable correspondence between the rankings for the ad hoc tasks in Section 4 and the rankings for the derived article retrieval measures. The correlation with the Focused task runs is much lower than with the Relevant in Context and Best in Context tasks. This makes sense, since the ranking of articles is an important part of the two “in context” tasks.

6.2 Article retrieval: Clicked pages

In addition to the topics created and assessed by INEX participants, we also included 150 queries derived from a proxy log, and can also construct pseudo-relevance judgments by regarding every clicked Wikipedia article as relevant.

Table 14. Top 10 Participants in the Ad Hoc Track: Article retrieval per task over judged topics (left) and clicked pages (right).

(a) Focused Task											
Participant	P5	P10	1/rank	map	bpref	Participant	P5	P10	1/rank	map	bpref
p92-manualQEIn*	0.6371	0.5843	0.8322	0.3601	0.3917	p5-Terrier	0.1594	0.0877	0.5904	0.5184	0.8266
p10-TOPXCOarti	0.5914	0.5386	0.8635	0.3489	0.3620	p6-inex08artB	0.1623	0.0870	0.5821	0.5140	0.8150
p5-GPX1COFOCp	0.5686	0.5214	0.7868	0.3390	0.3580	p92-autoindri0	0.1565	0.0884	0.5601	0.4853	0.8211
p37-kulcoeleme	0.5229	0.4500	0.7468	0.3240	0.3335	p60-JMUexpe142	0.1536	0.0862	0.5624	0.4853	0.8250
p78-FOER	0.5714	0.4986	0.7995	0.3230	0.3273	p48-LIGMLFOCRI	0.1449	0.0833	0.5191	0.4596	0.7153
p60-JMUexpe136	0.5429	0.4814	0.7843	0.3173	0.3380	p10-TOPXCOarti	0.1522	0.0841	0.5164	0.4538	0.8167
p25-weightedfi	0.4914	0.4600	0.7192	0.3164	0.3319	p78-FOER	0.1304	0.0819	0.4979	0.4404	0.8136
p29-VSMfbElts0	0.5486	0.4800	0.7955	0.3163	0.3385	p40-xfirmcos07	0.1217	0.0717	0.4301	0.3748	0.7184
p9-UHelRun293	0.5714	0.4957	0.7766	0.3113	0.3317	p55-KikoriFocu	0.1261	0.0732	0.4334	0.3727	0.7785
p6-inex08artB	0.5486	0.4757	0.7851	0.2990	0.3104	p22-EMSEFocus*	0.1203	0.0783	0.4233	0.3704	0.8105

(b) Relevant in Context Task											
Participant	P5	P10	1/rank	map	bpref	Participant	P5	P10	1/rank	map	bpref
p92-manualQEIn*	0.6371	0.5843	0.8322	0.3601	0.3917	p5-Terrier	0.1594	0.0877	0.5904	0.5184	0.8266
p5-GPX1CORICp	0.5686	0.5214	0.7868	0.3390	0.3580	p6-inex08artB	0.1623	0.0870	0.5821	0.5140	0.8150
p78-RICBest	0.5800	0.4943	0.8161	0.3371	0.3418	p60-JMUexpe150	0.1536	0.0862	0.5624	0.4853	0.8167
p60-JMUexpe150	0.5857	0.4886	0.8266	0.3107	0.3181	p92-autoindri0	0.1565	0.0884	0.5601	0.4853	0.8211
p10-TOPXCOallA	0.5257	0.4757	0.8226	0.3094	0.3275	p48-LIGMLRIC4O	0.1464	0.0841	0.5238	0.4647	0.7081
p4-SWKL200	0.5657	0.4943	0.7950	0.3086	0.3292	p78-RICBest	0.1348	0.0812	0.4979	0.4422	0.8126
p6-inex08artB	0.5486	0.4757	0.7851	0.2989	0.3104	p10-TOPXCOallA	0.1333	0.0775	0.5139	0.4397	0.7863
p56-VSMRIP05	0.5486	0.4514	0.7752	0.2869	0.3041	p72-UMDRic2	0.1275	0.0717	0.4560	0.4088	0.7526
p72-UMDRic2	0.5971	0.5157	0.8508	0.2715	0.3042	p4-SWKL200	0.1159	0.0732	0.4168	0.3701	0.8007
p22-EMSERICStr*	0.5029	0.4529	0.7079	0.2712	0.3054	p55-KikoriRele	0.1232	0.0710	0.4125	0.3501	0.7712

(c) Best in Context Task											
Participant	P5	P10	1/rank	map	bpref	Participant	P5	P10	1/rank	map	bpref
p78-BICER	0.6200	0.5257	0.8711	0.3753	0.3693	p5-Terrier	0.1594	0.0877	0.5904	0.5184	0.8266
p92-manualQEIn*	0.6371	0.5843	0.8322	0.3601	0.3917	p6-submitinex	0.1594	0.0862	0.5673	0.4976	0.8164
p5-GPX1COBICe	0.5686	0.5214	0.7868	0.3390	0.3580	p92-autoindri0	0.1565	0.0884	0.5601	0.4853	0.8211
p10-TOPXCOallB	0.5257	0.4757	0.8226	0.3261	0.3340	p60-JMUexpe151	0.1536	0.0855	0.5624	0.4844	0.8214
p25-weightedfi	0.4914	0.4600	0.7192	0.3224	0.3350	p78-BICPRplus	0.1522	0.0841	0.5432	0.4673	0.7799
p60-JMUexpe151	0.5857	0.4886	0.8266	0.3086	0.3178	p10-TOPXCOallB	0.1333	0.0775	0.5139	0.4398	0.8205
p6-submitinex	0.5457	0.4729	0.7793	0.2965	0.3081	p72-UMDBIC1	0.1275	0.0710	0.4482	0.4011	0.7398
p56-VSMRIP08	0.5486	0.4514	0.7752	0.2869	0.3041	p40-xfirmcos07	0.1217	0.0717	0.4301	0.3748	0.7160
p72-UMDBIC2	0.5886	0.5129	0.8440	0.2738	0.3016	p55-KikoriBest	0.1261	0.0732	0.4334	0.3727	0.7785
p12-p8u3exp501	0.4771	0.4343	0.6997	0.2709	0.3058	p56-VSMRIP08	0.1130	0.0659	0.3943	0.3445	0.7258

Table 15 shows the best run of the top 10 participating groups. The first column gives the participant, see Table 5 for the full name of group. The second and third column give the precision at ranks 5 and 10, respectively. The fourth column gives the mean reciprocal rank. The fifth column gives mean average precision. The sixth column gives binary preference measures (using the top R judged non-relevant documents). Compared to the judged topics, we immediately see much lower scores for the early precision measures (precision at

Table 15. Top 10 Participants in the Ad Hoc Track: Clicked articles.

Participant	P5	P10	1/rank	map	bpref
p5-Terrier	0.1594	0.0877	0.5904	0.5184	0.8266
p6-inex08artB	0.1623	0.0870	0.5821	0.5140	0.8150
p60-JMUexpe150	0.1536	0.0862	0.5624	0.4853	0.8167
p92-autoindri0	0.1565	0.0884	0.5601	0.4853	0.8211
p78-BICPRplus	0.1522	0.0841	0.5432	0.4673	0.7799
p48-LIGMLRIC4O	0.1464	0.0841	0.5238	0.4647	0.7081
p10-TOPXCOarti	0.1522	0.0841	0.5164	0.4538	0.8167
p72-UMDRic2	0.1275	0.0717	0.4560	0.4088	0.7526
p40-xfirmcos07	0.1217	0.0717	0.4301	0.3748	0.7184
p55-KikoriFocu	0.1261	0.0732	0.4334	0.3727	0.7785

5 and 10, and reciprocal ranks), while at the same time higher scores for the overall measures (map and bpref). This is a result of the very low numbers of relevant documents, 1.8 on average, that make it impossible to get a grips on recall aspects. The runs ranked first (*p5-Terrier*), fourth (*p92-autoindri0*), and seventh (*p10-TOPXCOarti*) retrieve exclusively full articles. Again, it is no great surprise that these runs do well for the task of article retrieval.

The resulting ranking is quite different from the article ranking based on the judged ad hoc topics in Table 13. They have only one run in common, although they agree on five of the ten participants. Looking, more formally, at the system rank correlations between the two types of article retrieval we see the following.

- Over all 163 submissions, the system rank correlation is 0.359.
- Over the 76 Focused task submissions, the correlation is 0.363.
- Over the 49 Relevant in task submissions, the correlation is 0.374.
- Over the 38 Best in Context task submissions, the correlation is 0.388.

Hence the judged topics above and the topics derived from the proxy log vary considerable. A large part of the explanation is the dramatic difference between the numbers of relevant articles, with 69.3 on average for the judged topics and 1.8 on average for the proxy log topics.

7 Discussion and Conclusions

In this paper we provided an overview of the INEX 2008 Ad Hoc Track that contained three tasks: For the *Focused Task* a ranked-list of non-overlapping results (elements or passages) was required. For the *Relevant in Context Task* non-overlapping results (elements or passages) grouped by the article that they belong to were required. For the *Best in Context Task* a single starting point (element’s starting tag or passage offset) per article was required. We discussed the results for the three tasks, and analysed the relative effectiveness of element and passage runs, and of keyword (CO) queries and structured queries (CAS). We also look at effectiveness in term of article retrieval, both using the judged topics and using queries and clicks derived from a proxy log.

When examining the relative effectiveness of CO and CAS we found that for all tasks the best scoring runs used the CO query. This is in contrast with earlier results showing that structural hints can help promote initial precision [8]. Part of the explanation may be in the low number of CAS submissions (28) in comparison with the number of CO submissions (108). Only 39 of the 70 judged topics had a non-trivial CAS query, and the majority of those CAS queries made only reference to particular tags and not on their structural relations. This may have diminished the value of the CAS query in comparison with earlier years.

Given the efforts put into the fair comparison of element and passage retrieval approaches, the number of passage and FOL submissions was disappointing. Twenty-two submissions used ranges of elements or FOL passage results, whereas 118 submissions used element results. In addition, many of the passage or FOL submissions used exclusively full articles as results. Although we received too few non-element runs to draw clear conclusions, we saw that the passage based approaches were competitive, but not superior to element based approaches. This outcome is consistent with earlier results using passage-based element retrieval [4, 5].

As in earlier years, we saw that article retrieval is a reasonably effective at XML-IR: for each of the ad hoc tasks there were three article-only runs among the best runs of the top 10 groups. When looking at the article rankings inherent in all Ad Hoc Track submissions, we saw that again three of the best runs of the top 10 groups in terms of article ranking (across all three tasks) were in fact article-only runs. This also suggests that element-level or passage-level evidence is still valuable for article retrieval. When comparing the system rankings in terms of article retrieval with the system rankings in terms of the ad hoc retrieval tasks, over the exact same topic set, we see a reasonable correlation especially for the two “in context” tasks. The systems with the best performance for the ad hoc tasks, also tend to have the best article rankings. When we look at a different topic set derived from a proxy log, and a shallow set of clicked pages rather than a full-blown IR test collection, we see notable differences. Given the low number of relevant articles (1.8 on average) compared to the ad hoc judgments (69.3 on average), the clicked pages focus exclusively on precision aspects. This leads to a different system ranking, although there is still some agreement on the best groups. The differences between these two sets of topics require further analysis.

Finally, the Ad Hoc Track had two main research questions. The first main research question was the comparative analysis of element and passage retrieval approaches, hoping to shed light on the value of the document structure as provided by the XML mark-up. We found that the best performing system used predominantly element results, although the number of non-element retrieval runs submitted is too low to draw any definite conclusions. The second main research question was to compare focused retrieval directly to traditional article retrieval. We found that the best scoring Ad Hoc Track submissions also tend to have the best article ranking, and that the best article rankings were generated using element-level evidence. For both main research questions, we hope and expect that the resulting test collection will prove its value in future use. After

all, the main aim of the INEX initiative is to create bench-mark test-collections for the evaluation of structured retrieval approaches.

Acknowledgments Jaap Kamps was supported by the Netherlands Organization for Scientific Research (NWO, grants 612.066.513, 639.072.601, and 640.001.501).

Bibliography

- [1] J. P. Callan. Passage-level evidence in document retrieval. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 302–310. Springer-Verlag, New York NY, 1994.
- [2] C. L. A. Clarke. Range results in XML retrieval. In *Proceedings of the INEX 2005 Workshop on Element Retrieval Methodology*, pages 4–5, Glasgow, UK, 2005.
- [3] L. Denoyer and P. Gallinari. The Wikipedia XML Corpus. *SIGIR Forum*, 40:64–69, 2006.
- [4] W. Huang, A. Trotman, and R. A. O’Keefe. Element retrieval using a passage retrieval approach. In *Proceedings of the 11th Australasian Document Computing Symposium (ADCS 2006)*, pages 80–83, 2006.
- [5] K. Y. Itakura and C. L. A. Clarke. From passages into elements in XML retrieval. In *Proceedings of the SIGIR 2007 Workshop on Focused Retrieval*, pages 17–22. University of Otago, Dunedin New Zealand, 2007.
- [6] J. Kamps and M. Koolen. On the relation between relevant passages and XML document structure. In *Proceedings of the SIGIR 2007 Workshop on Focused Retrieval*, pages 28–32. University of Otago, Dunedin New Zealand, 2007.
- [7] J. Kamps, M. Marx, M. de Rijke, and B. Sigurbjörnsson. The importance of morphological normalization for XML retrieval. In *Proceedings of the First INEX Workshop*, pages 41–48. ERCIM, 2003.
- [8] J. Kamps, M. Marx, M. de Rijke, and B. Sigurbjörnsson. Articulating information needs in XML query languages. *Transactions on Information Systems*, 24:407–436, 2006.
- [9] J. Kekäläinen and K. Järvelin. Using graded relevance assessments in IR evaluation. *Journal of the American Society for Information Science and Technology*, 53:1120–1129, 2002.
- [10] J. A. Thom and J. Pehcevski. How well does best in context reflect ad hoc XML retrieval. In *Pre-Proceedings of INEX 2007*, pages 124–125, 2007.
- [11] A. Trotman and S. Geva. Passage retrieval and other XML-retrieval tasks. In *Proceedings of the SIGIR 2006 Workshop on XML Element Retrieval Methodology*, pages 43–50. University of Otago, Dunedin New Zealand, 2006.

A Appendix: Full run names

Group	Run	Label	Task	Query	Results	Notes
4	151	p4-SWKL200	RiC	CO	Pas	
4	152	p4-WHOLEDOC	RiC	CO	Ele	Article-only
4	153	p4-WHOLEDOCPA	RiC	CO	Pas	Article-only
5	122	p5-Terrier	BiC	CO	Pas	Article-only
5	123	p5-Terrier	Foc	CO	Pas	Article-only
5	124	p5-Terrier	RiC	CO	Pas	Article-only
5	133	p5-GPX2COFOCP	Foc	CO	Pas	Article-only
5	138	p5-GPX1COBICe	BiC	CO	Ele	
5	139	p5-GPX1COFOCe	Foc	CO	Ele	
5	140	p5-GPX1CORICe	RiC	CO	Ele	
5	141	p5-GPX3COSBIC	BiC	CAS	Ele	
5	142	p5-GPX3COSFOC	Foc	CAS	Ele	
5	143	p5-GPX3COSRIC	RiC	CAS	Ele	
5	144	p5-GPX1COBICp	BiC	CO	Pas	Article-only
5	145	p5-GPX1COFOCP	Foc	CO	Pas	Article-only
5	146	p5-GPX1CORICp	RiC	CO	Pas	Article-only
6	255	p6-submitinex	BiC	CO	FOL	Article-only
6	264	p6-inex08artB	RiC	CAS	Ele	
6	265	p6-inex08artB	RiC	CO	Ele	
6	269	p6-inex08artB	RiC	CO	Ele	
6	270	p6-inex08artB	Foc	CAS	Ele	
6	271	p6-inex08artB	Foc	CO	Ele	
6	274	p6-inex08artB	Foc	CO	Ele	
6	276	p6-inex08artB	Foc	CO	Ele	
9	174	p9-UHelRun293	Foc	CO	Ele	
9	176	p9-UHelRun394	Foc	CO	Ele	
10	91	p10-TOPXCOallF	Foc	CO	Ele	
10	92	p10-TOPXCOallB	BiC	CO	Ele	
10	93	p10-TOPXCOallA	RiC	CO	Ele	
10	207	p10-TOPXCOarti	Foc	?	Ele	Article-only
12	97	p12-p8u3exp501	BiC	CO	Ele	
12	100	p12-p8u3exp511	RiC	CO	Ele	
14	205	p14-T2FBCOPARA	Foc	CO	Ele	
16	233	p16-009RunofUn	RiC	CAS	Ele	
16	234	p16-006RunofUn	BiC	CAS	Ele	
16	244	p16-001RunofUn	Foc	CAS	Ele	
22	62	p22-EMSEFocus	Foc	CO	Ele	Manual Invalid
22	66	p22-EMSEFocus	Foc	CO	FOL	Manual
22	68	p22-EMSERICStr	RiC	CO	Ele	Manual Invalid
25	30	p25-RUCLLP08	Foc	CAS	Ele	
25	278	p25-weightedfi	Foc	CO	Ele	

Continued on Next Page...

Group	Run	Label	Task	Query	Results	Notes
25	282	p25-weightedfi	BiC	CO	Ele	
29	238	p29-VSMfbElts0	Foc	CO	Ele	
29	253	p29-LMnofb020	Foc	CO	Ele	Article-only
37	227	p37-kulcaselem	Foc	CAS	Ele	
37	230	p37-kulcoeleme	Foc	CO	Ele	
40	54	p40-xfirmbicco	BiC	CO	Ele	
40	296	p40-xfirmcos07	BiC	CAS	Ele	
40	297	p40-xfirmcos07	Foc	CAS	Ele	Invalid
42	299	p42-B2U0visith	Foc	CAS	Ele	
48	59	p48-LIGMLFOCRI	Foc	CO	Ele	
48	72	p48-LIGMLRIC4O	RiC	CO	Ele	
55	279	p55-KikoriFocu	Foc	CAS	Ele	Invalid
55	280	p55-KikoriRele	RiC	CAS	Ele	Invalid
55	281	p55-KikoriBest	BiC	CAS	Ele	
56	190	p56-VSMRIP02	Foc	CAS	Ele	
56	197	p56-VSMRIP04	RiC	CO	Ele	Article-only
56	199	p56-VSMRIP05	RiC	CAS	Ele	Article-only
56	202	p56-VSMRIP08	BiC	CAS	Ele	
56	224	p56-VSMRIP09	BiC	CO	Ele	
60	11	p60-JMUexpe136	Foc	CO	Ele	Article-only
60	53	p60-JMUexpe142	Foc	CO	Ele	
60	81	p60-JMUexpe150	RiC	CO	Ele	Invalid
60	82	p60-JMUexpe151	BiC	CO	Ele	Invalid
72	106	p72-UMDFocused	Foc	CO	Ele	
72	154	p72-UMDBIC1	BiC	CO	Ele	
72	155	p72-UMDBIC2	BiC	CO	Ele	
72	277	p72-UMDRic2	RiC	CO	Ele	
78	156	p78-FOER	Foc	CO	Ele	
78	157	p78-FOERStep	Foc	CO	Ele	
78	160	p78-BICER	BiC	CO	Ele	
78	163	p78-BICPRplus	BiC	CO	Pas	
78	164	p78-RICBest	RiC	CO	Ele	
92	177	p92-autoindri0	BiC	CO	Ele	Article-only
92	178	p92-autoindri0	Foc	CO	Ele	Article-only
92	179	p92-autoindri0	RiC	CO	Ele	Article-only
92	183	p92-manualQEin	BiC	CO	Ele	Manual Article-only
92	184	p92-manualQEin	Foc	CO	Ele	Manual Article-only
92	185	p92-manualQEin	RiC	CO	Ele	Manual Article-only