

# Overview of the INEX 2012 Linked Data Track

Qiuyue Wang<sup>1</sup>, Jaap Kamps<sup>2</sup>, Georgina Ramírez Camps<sup>3</sup>, Maarten Marx<sup>2</sup>,  
Anne Schuth<sup>2</sup>, Martin Theobald<sup>4</sup>, Sairam Gurajada<sup>4</sup>, and Arunav Mishra<sup>4</sup>

<sup>1</sup>Renmin University of China, Beijing, China

<sup>2</sup>University of Amsterdam, Amsterdam, The Netherlands

<sup>3</sup>Universitat Pompeu Fabra, Barcelona, Spain

<sup>4</sup>Max Planck Institute for Informatics, Saarbrücken, Germany

**Abstract.** This paper provides an overview of the Linked Data Track that was newly introduced to the set of INEX tracks in 2012.

## 1 Introduction

The goal of the new Linked Data Track was to investigate retrieval techniques over a combination of textual and highly structured data, where rich textual contents from Wikipedia articles serve as the basis for retrieval and ranking, while additional RDF properties carry key information about semantic relations among entities that cannot be captured by keywords alone. Our intension in organizing this new track thus follows one of the key themes of INEX, namely to explore and investigate if and how structural information could be exploited to improve the effectiveness of *ad-hoc* retrieval. In particular, we were interested in how this combination of data could be used together with structured queries to help users navigate or explore large sets of results (a task that is well-known from *faceted search* systems), or to address Jeopardy-style natural-language clues and questions (known, for example, from recent *question answering* settings over linked data collections, see for example [6]). The Linked Data Track thus aims to close the gap between IR-style keyword search and semantic-web-style reasoning techniques, with the goal to bring together different communities and to foster research at the intersection of Information Retrieval, Databases, and the Semantic Web.

As its core collection, the Linked Data Track employs a fusion of XML-ified Wikipedia articles with RDF properties from both DBpedia [4] and YAGO2 [5], the latter of which contain the article entity as either their subject (first argument) or object (second argument). The core data collection was based on the popular MediaWiki format<sup>1</sup>, where we additionally replaced all Wiki-markup by syntactically valid XML tags, attributes, and CDATA sections. In addition, all internal Wikipedia links (including the article entity itself) have been enriched with links to both their corresponding DBpedia and YAGO2 entities (as far as available). In addition, participants were explicitly encouraged to make use of

<sup>1</sup> <http://dumps.wikimedia.org/enwiki/20110722/>

more RDF facts available from DBpedia and YAGO2, in particular for processing the reasoning-related faceted search and Jeopardy topics. For INEX 2012, we explored three different retrieval tasks:

- The classic **Ad-hoc Retrieval Task** investigates informational queries to be answered mainly by the textual contents of the Wikipedia articles.
- The **Faceted Search Task** employs a hand-crafted hierarchy of facets and facet-values obtained from DBpedia that aim to guide the searcher toward relevant information.
- The new **Jeopardy Task** employs natural-language Jeopardy clues which are manually translated into a semi-structured query format based on SPARQL with keyword filter conditions.

## 2 Data Collection

The new Wikipedia-LOD (v1.1) collection is hosted by the Max Planck Institute for Informatics and has been made available for download in May 2012 from the following link: <http://www.mpi-inf.mpg.de/inex-lod/wikipedia-lod-2012/>

The collection consists of 3 compressed tar.gz files and contains an overall amount of 3.1 Million individual XML articles. The uncompressed size of the collection is 61 GB. A detailed DTD file that describes the structure of the XML collection is also available from the above URL. Each Wikipedia-LOD article consists of a mixture of XML tags, attributes, and CDATA sections, containing infobox attributes, free-text contents, describing the entity or category that the article captures, and a section with both DBpedia and YAGO2 properties that are related to the article’s entity. All sections contain links to other Wikipedia articles (including links to the corresponding DBpedia and YAGO2 resources), Wikipedia categories, and external Web pages.

Figure 1 shows an example of an XML-ified Wikipedia article about the entity `Albert_Einstein` by depicting the two main sections of the article:

- i) the Wikipedia section, containing an XML-ified infobox, enhanced links pointing to DBpedia and YAGO2, and Wikipedia text contents with more XML markup, and
- ii) the Linked Data section with RDF triples imported from both DBpedia and YAGO2 that contain the entity `Albert_Einstein` as either their subject or object.

**Wikipedia To WikiXML Parser.** For converting the raw Wikipedia articles into our XML format, we used a parser derived from the wiki2xml parser [3] provided by MediaWiki [1]. The parser generates an XML file from the raw Wikipedia article (originally in Wiki markup) by transforming infobox information to a proper XML representation, comprehending links with DBpedia and YAGO2 entities, and finally annotating each article with a list of RDF properties from the DBpedia and YAGO2 knowledge sources.

```

<!DOCTYPE html>
<loidxml xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance' xml:lang='en'
xmlns:xhtml='http://www.w3.org/1999/xhtml' encoding='UTF-8'>
<article title='Albert_Einstein'>
<wikipedia>
<template type='Metadata'>
<tag name='id'>736</tag>
<tag name='title'>Albert_Einstein</tag>
</template>
<infobox type='person'>
<tag name='name'>Albert Einstein</tag>
<tag name='image'>Einstein 1921 portrait2.jpg</tag>
.....
</infobox>
.....
<paragraph>Albert Einstein was born in<space/>
.....
<link>
<wikilink href='http://en.wikipedia.org/wiki/Ulm'>Ulm</wikilink>
<dbpedia href='http://dbpedia.org/resource/Ulm' />
<yago ref='Ulm' />
</link>
.....
</wikipedia>
<dbpediaproperties>
<property name='http://dbpedia.org/property/birthPlace'>
<object name='http://dbpedia.org/resource/German_Empire' />
</property>
.....
</dbpediaproperties>
<yagoproperties>
<property name='isKnownFor'><object name='Einstein_field_equations' /></property>
<property name='type'><object name='wikicategory_German_pacifists' /></property>
<property name='type'><object name='wikicategory_German_philosophers' /></property>
<property name='type'><object name='wikicategory_German_physicists' /></property>
</yagoproperties>
</article>
</loidxml>

```

**Fig. 1.** XML-ified Wikipedia articles with DBpedia and YAGO2 properties

**Collection Statistics.** The Wikipedia-LOD collection currently contains 3.1 Million XML documents in 3 compressed tar.gz files counting to the size of 61 GB in uncompressed form. Table 1 provides more detailed numbers about different properties of the collection.

**Linked Data Sources.** In addition to the new core collection, which is based on XML-ified Wikipedia articles, the Linked Data Track explicitly encourages (but does not require) the use of current Linked Open Data dumps for DBpedia (v3.7) and YAGO2, which are available from the following URLs:

- DBpedia v3.7 (created in July 2011):  
<http://downloads.dbpedia.org/3.7/en/>
- YAGO2 core and full dumps (created on 2012-01-09):  
<http://www.mpi-inf.mpg.de/YAGO2-naga/YAGO2/>

Property	Count
XML Documents	3,164,041
XML Elements	1,173,255,397
Wikipedia Category Articles	266,134
Wikipedia Entity Articles	2,053,050
Wikipedia Entity Articles with Infoboxes	907,304
Other Wikipedia Articles	844,857
Resolved DBpedia Links	36,941,795
Resolved YAGO2 Links	32,941,667
Intra-Wiki Links	22,235,753
External Web Links	7,214,827
Imported DBpedia Properties	168,374,863
Imported YAGO2 Properties	23,634,511

**Table 1.** Wikipedia-LOD (v1.1) Collection Statistics

DBpedia and YAGO2 are two comprehensive, common-sense knowledge bases providing structured information that has been semi-automatically extracted mostly from Wikipedia infoboxes and categories. Both knowledge bases focus on extracting attribute-value pairs from Wikipedia infoboxes and category lists, which serve as basis for applying various information extraction techniques. They also contain geo-coordinates, links between Wikipedia pages, redirection and disambiguation pages, external links, and much more. Each Wikipedia page corresponds to a resource in DBpedia and YAGO2. The connection between the data sets is given in the "wikipedia\_links\_en.nt" file from DBpedia. The following entry, for example,

```
<http://dbpedia.org/resource/AccessibleComputing>
<http://xmlns.com/foaf/0.1/page>
<http://en.wikipedia.org/wiki/AccessibleComputing>
```

connects the DBpedia entity with the URI <http://dbpedia.org/resource/AccessibleComputing> with the Wikipedia page that is available under the URI <http://en.wikipedia.org/wiki/AccessibleComputing>.

The Linked Data Track was explicitly intended to be an "open track" and thus invited participants to include more Linked Data sources (see, for example, <http://linkeddata.org>) or other sources that go beyond "just" DBpedia and YAGO2. Any inclusion of further data sources was welcome, however, workshop submissions and follow-up research papers should explicitly mention these sources when describing their approaches.

### 3 Retrieval Tasks and Topics

#### 3.1 Ad-hoc Task and Faceted Search Tasks

The Ad-hoc Task is to return a ranked list of results (Wikipedia pages) estimated relevant to the user's information need, which is typically formulated into a

keyword query. Given an exploratory or broad query, the search system may return a large number of results. Faceted search is a way to help users navigate through the large set of results to quickly identify the results of interest. It presents the user a list of facet-values to refine the query. After the user choosing from the suggested facet-values, the result list is narrowed down and then the system may present a new list of facet-values for the user to further refine the query. The interactive process continues until the user finds the items of interest. One of the key issues in faceted search systems is to recommend appropriate facet-values to help the user quickly identify what he/she really wants in the large set of results. The task aims to investigate different techniques of recommending facet-values.

This year, we did not ask participants to submit ad-hoc or faceted search topics. We generated and collected the topics from the following three sources. Firstly, we built a three-level hierarchy of topics as described in [7]. For example,

```
Vietnam
  Vietnam war
  Vietnam war movies
  Vietnam war facts
  Vietnam food
    Vietnam food recipes
    Vietnam food blog
  Vietnam travel
    Vietnam travel national park
    Vietnam travel airports
```

The topics on the top level are general topics, e.g., “Vietnam”. We randomly created 5 general topics, i.e. “Vietnam”, “guitar”, “tango”, “bicycle”, and “music”. For each general topic, we typed it into Google, and from Google’s online suggestions, we chose 3 subtopics. For example, when you type in “Vietnam”, Google may suggest “Vietnam war”, “Vietnam food” or “Vietnam travel”, and so on, which can be viewed as subtopics to “Vietnam”. Furthermore, for each subtopic, we selected 2 sub-subtopics using Google Suggest again. Thus we formed a three-level hierarchy of topics, with 5 general topics, 15 subtopics and 30 sub-subtopics. Since the relevant answers for a topic can be treated as the union of the relevant answers of all its subtopics, only the leaf-level topics, i.e. 30 sub-subtopics need to be assessed. So we put the 30 sub-subtopics to the Ad-hoc Task and 20 non-leaf level topics to the Faceted Search Task. The relevance results for the ad-hoc topics will serve as the relevant results to their corresponding faceted search topics.

Secondly, we selected 20 topics from INEX 2009 and 2010 Ad-hoc Tracks to compare the performance of different data collections. Since we want to select challenging topics, we took 40 worst performed topics (with lowest average precisions) from the INEX 2009 Ad-hoc Track and 30 worst performed topics from the INEX 2010 Ad-hoc Track, and then randomly selected 10 topics from each set. In this process, we also found some natural general topics, “Normandy”, “museum” and “social networ”, which have multiple subtopics among the 20

topics that we collected. So we added the 3 topics to the set of faceted search topics.

Thirdly, to compare the performance of structured queries that were used in Jeopardy Task and unstructured queries, we added all the 90 keyword titles of Jeopardy topics into the set of ad-hoc topics. In total, we collected 140 ad-hoc topics and 23 faceted search topics, which are in the same format as that in previous years [8].

### 3.2 Jeopardy Task

The new Jeopardy Task investigated retrieval techniques over a set of 90 natural-language Jeopardy-style clues and questions, which have been manually translated into SPARQL query patterns that were enhanced with keyword-based filter conditions. Specifically, we investigated a data model, where every entity (in DBpedia or YAGO2) is associated with the Wikipedia article (contained in the Wikipedia-LOD v1.1 collection) that describes this entity. An XML file with 90 Jeopardy-style topics was made available available for download in June 2012 under the following URL:

<http://www.mpi-inf.mpg.de/inex-lod/LDT-2012-jeopardy-topics.xml>

For example, topic no. 2012301 from the current set of Jeopardy topics looks as follows:

```
<topic id="2012301" category="LAKES">
  <jeopardy_clue>Niagara Falls has its source of origin
  from this lake. </jeopardy_clue>
  <keyword_title>Niagara Falls source lake</keyword_title>
  <sparql_ft>
    Select ?q Where {
      <http://dbpedia.org/resource/Niagara_Falls>
        <http://dbpedia.org/property/watercourse> ?o .
      ?o <http://dbpedia.org/ontology/origin> ?q .
      Filter FTContains(?o, "river water course niagara") .
      Filter FTContains(?q, "lake origin of")}
    </sparql_ft>
</topic>
```

The `<jeopardy_clue>` element contains the original Jeopardy clue as a natural-language sentence; the `<keyword_title>` element contains a set of keywords that have been manually extracted from this title and will be reused as part of the Ad-hoc Retrieval Task; and the `<sparql_ft>` element contains a formulation of the natural-language sentence into a corresponding SPARQL pattern. The `<category>` attribute of the `<topic>` element may be used as an additional hint for disambiguating the query.

In the above query, the DBpedia entity `http://dbpedia.org/resource/Niagara_Falls` has been marked as the subject of the first triplet pattern, while both the object of the first triplet pattern and the subject and object of the second

triplet pattern are unknown. The two `FTContains` filter conditions however restrict both these subjects and objects to entities that should be associated with the keywords “river water course niagara” and “lake origin” via the content of their corresponding Wikipedia articles, respectively. The result of this query is exactly one target entity, namely the DBpedia resource `http://dbpedia.org/resource/Lake.Erie`.

Since this particular variant of processing SPARQL queries with full-text filter conditions is not a default functionality of current SPARQL engines (and queries should not be run against a standard RDF collection such as DBpedia or YAGO2 alone), participants were encouraged to develop individual solutions to index both the RDF and textual contents of the Wikipedia-LOD collection in order to process these queries. Adding full-text search to SPARQL queries is an ongoing research issue. While initial implementations and syntax proposals exist (see for example [2]), we are not aware of any SPARQL engine that currently allows for associating and indexing entire text documents along with RDF resources. We also remark that this particular LOD data model differs from most current SPARQL full-text approaches, as we impose keyword conditions over individual entities (resources) rather than entire facts (triplets).

## 4 Run Submissions

All run submissions were to be uploaded via the INEX website via the URL: `https://inex.mmci.uni-saarland.de/`. The due date for the submission of all LOD runs was July 14, 2012.

### 4.1 Ad-Hoc and Jeopardy Tasks

For the Ad-hoc and Jeopardy Tasks, each run must contain a maximum of 1,000 results per topic, ordered by decreasing value of relevance. For the Ad-hoc Task, each result is a Wikipedia article uniquely identified by its page ID. For the Jeopardy Task however, each query result could be a set of entities (identified by their corresponding Wikipedia page IDs) in case that the select clause contains more than one query variables. For relevance assessment and evaluation of the results, we require submission files to be in the familiar TREC format, with each row representing a single query result. In case the select clause contains more than one query variable as in a Jeopardy topic, the row should consist of a comma- or semicolon-separated list of target entity ID’s. This list of entities must reflect the order of query variables as specified by the select clause of the Jeopardy topic.

```
<qid> Q0 <page_id_list> <rank> <rsv> <run_id>
```

Where:

- The first column is the topic number.

- The second column is the query number within that topic. This is currently unused and should always be Q0.
- The third column is a comma- or semicolon-separated list the ID’s of the resulting Wikipedia page(s).
- The fourth column is the rank of the result.
- The fifth column shows the score (integer or floating point) that generated the ranking.
- The sixth column is called the “run tag” and should be a unique identifier for your group AND for the method used. Run tags must contain 12 or fewer letters and numbers, with NO punctuation, to facilitate labeling graphs with the tags.

An example submission thus may look as follows:

```
2012301 Q0 12 1 0.9999 2012UniXRun1
2012301 Q0 997 2 0.9998 2012UniXRun1
2012301 Q0 9989 3 0.9997 2012UniXRun1
```

Here we have three results for topic “2012301”. The first result is the entity (i.e. Wikipedia page) with ID “12”. The second result is the entity with ID “997”, and the third result is the entity with ID “9989”.

## 4.2 Faceted Search Task

For the Faceted Search Task, the organizers will provide a result file, which contains a result list of maximum 2000 results for each general topic. Based on the reference result file, a run submitted by a participant should be a XML file conforming to the following DTD, which contains a hierarchy of recommended facet-values for each topic, in which each node represents a facet-value and all of its children constitute the newly recommended facet-value list when the user selects this facet-value to refine the query. The maximum fan-out of each node in the hierarchy is restricted to be 20.

```
<!ELEMENT run (topic+)>
<!ATTLIST run rid ID #REQUIRED>
<!ELEMENT topic (fv+)>
<!ATTLIST topic tid ID #REQUIRED>
<!ELEMENT fv (fv*)>
<!ATTLIST fv f CDATA #REQUIRED
           v CDATA #REQUIRED>
```

Where:

- The root element is `<run>`, which has an ID type attribute, `rid`, representing the unique identifier of the run.
- The `<run>` contains one or more `<topic>`’s. The ID type attribute, `tid`, in each `<topic>` gives the topic number.



- Each <topic> has a hierarchy of <fv>’s. Each <fv> shows a facet-value pair, with f attribute being the facet and v attribute being the value. All the possible facet-value pairs are from the triples in DBpedia or YAGO2.
- The <fv>’s can be nested to form a hierarchy of facet-values.

An example submission is:

```

<run rid=2012UniXRun1>
<topic tid=2012001>
<fv f=dbpedia-owl:date v=1955-11-01>
  <fv f=dbpedia-owl:place v=dbpedia:South_Vietnam>
    <fv f=rdf:type v=dbpedia-owl:MilitaryConflict/>
    <fv f=rdf:type v=dbpedia-owl:Country/>
  </fv>
  <fv f=dbpedia-owl:place v=dbpedia:North_Vietnam>
    <fv f=rdbpprob:capital v=dbpedia:Ho_Chi_Minh_City/>
  </fv>
</fv>
...
</topic>
<topic tid=2012002>
...
</topic>
...
</run>

```

Here for the topic “2012001”, the faceted search system first recommends the facet-value condition “dbpedia-owl:date = 1955-11-01” among other facet-value conditions, which are its siblings. If the user selects this condition to refine the query, the system will recommend a new list of facet-value conditions, which are “dbpedia-owl:place = dbpedia:South\_Vietnam” and “dbpedia-owl:place = dbpedia:North\_Vietnam”. If the user then selects “dbpedia-owl:plac = dbpedia:North\_Vietnam”, the system will recommend the facet-value condition “rdbprob:capital = dbpedia:Ho\_Chi\_Minh\_City”. Note that no facet-value condition may occur twice on a path in the hierarchy.

## 5 Relevance Assessments and Evaluation Metrics

In total 20 ad-hoc search runs were submitted by 7 participants, i.e., *Ecole des Mines de Saint-Etienne* (EMSE), *Kasetsart University*, *Renmin University of China*, *University of Otago*, *Oslo University College*, *University of Amsterdam*, *Norwegian University of Science and Technology* (NTNU), and 5 valid Jeopardy runs were submitted by 2 participants, i.e., *Kasetsart University* and *Max-Planck Institute for Informatics* (MPI).

Assessment was done using the Amazon Mechanical Turk. We did not assess the 20 topics from the INEX 2009 and 2010 Ad-hoc Tracks as we could use the

assessment results done in previous years. We assessed the 30 sub-subtopics and 50 Jeopardy topics randomly selected from the 90 ones. For each sub-subtopic, we pooled all the submitted runs in a round-robin manner, and then picked up the top 200 results to be assessed. For each selected Jeopardy topic, we pooled the results in the same way and picked up the top 100 results to be assessed as in general Jeopardy Task can be viewed a known-item search.

The TREC MAP metric, as well as P@5, P@10, P@20 and so on, was used to measure the performance of all ad-hoc and Jeopardy runs. For the Faceted Search Task, we use the same metrics as that used in last year [?] to evaluate the runs.

## 6 Results

### 6.1 Ad-hoc and Jeopardy Task Results

As mentioned above, 140 ad-hoc topics were collected from three different sources: sub-subtopics, old topics from INEX 2009 and 2010, and keyword titles of Jeopardy topics. Among them, the 30 sub-subtopics, 20 old topics and 50 Jeopardy topics have assessment results. In this section, we will first present the evaluation results over the whole set of ad-hoc topics for all the submitted runs, and then analyze the effectiveness of the runs for each of the three sets of topics.

There are 20 runs submitted to the Ad-hoc Task by 7 participating groups. For each group, we selected its best performing run in terms of MAP, since MAP averages reasonably well over all topic types. Table 2 shows an overview of the 7 best performing runs from different groups. Over all topics, the best scoring run is from the Renmin University of China with a MAP of 0.2776 and also highest 1/rank, P@5, P@10, P@20 and P@30. Second best scoring team is University of Otago (0.2721). Third best scoring team is Ecole des Mines de Saint-Etienne (0.2609). Interpolated precision against recall is plotted in Fig 2, which shows little differences among the 3-4 best performing runs. The best performing runs are quite similar actually.

Table 3 shows the results over the 30 sub-subtopics. Since University of Amsterdam did not submit any results on sub-subtopics, there are only 6 instead of 7 runs in the table. We see that Renmin University of China (0.33365), University of Otago (0.3081), and Ecole des Mines de Saint-Etienne (0.2991) are still the 3 best performing groups.

Table 4 shows the results over the 20 old topics from INEX 2009 and 2010 Ad-hoc Tracks, now again evaluated by MAP. There are only 6 runs in the table since Oslo University College did not submit any results on this set of topics. We see that Renmin University of China still performs the best in terms of MAP (0.0936), and University of Amsterdam runs the second with the best 1/rank and P@5. The MAPs are commonly very low for this set of topics. This is no surprise since these are “hard” topics from previous years.

Table 5 shows the results over only the Jeopardy topics, now evaluated by the mean reciprocal rank (1/rank). There are 7 groups submitted results to the

Run	MAP	1/rank	P@5	P@10	P@20	P@30
Renmin-LDT2012_adhoc_ruc_comb07	<b>0.2776</b>	<b>0.7778</b>	<b>0.452</b>	<b>0.389</b>	<b>0.3235</b>	<b>0.2823</b>
Otago-ou2012pr09	0.2721	0.745	0.444	0.382	0.323	0.279
EMSE-run-085	0.2609	0.7131	0.444	0.367	0.3055	0.2663
NTNU-run1	0.2459	0.7145	0.436	0.372	0.3015	0.255
Amsterdam-inex12LDT.adhoc.baseline.LM	0.2187	0.7481	0.3829	0.2929	0.2114	0.1729
Kasetsart-kas16-PHR	0.1074	0.6718	0.3783	0.313	0.2489	0.2152
Oslo-result.fl	0.0046	0.037	0	0	0	0.0333

**Table 2.** Best performing runs (only showing one run per group) based on MAP over all the assessed ad-hoc topics.

Run	MAP	1/rank	P@5	P@10	P@20	P@30
Renmin-LDT2012_adhoc_ruc_comb15	<b>0.3365</b>	<b>0.8511</b>	<b>0.6067</b>	<b>0.6</b>	<b>0.5617</b>	<b>0.5167</b>
Otago-ou2012pr09	0.3081	0.8522	0.62	0.58	0.5467	0.4989
EMSE-run-086	0.2991	0.7356	0.58	0.5667	0.535	0.5067
NTNU-run1	0.2693	0.8122	0.6	0.5533	0.505	0.46
Kasetsart-kas16-EXT	0.1312	0.543	0.3154	0.3231	0.3173	0.3013
Oslo-result.fl	0.0046	0.037	0	0	0	0.0333


**Table 3.** Best performing runs (only showing one run per group) based on MAP over the 30 sub-subtopics

Run	MAP	1/rank	P@5	P@10	P@20	P@30
Renmin-LDT2012_adhoc_ruc_comb1	<b>0.0936</b>	<b>0.6845</b>	<b>0.33</b>	<b>0.29</b>	<b>0.2225</b>	<b>0.195</b>
Amsterdam-inex12LDT.adhoc.baseline.LM	0.0895	0.7146	0.34	0.29	0.22	0.1867
Otago-ou2012pr10	0.0836	0.5717	0.31	0.26	0.205	0.1783
EMSE-run-085	0.0782	0.5916	0.3	0.225	0.1875	0.1633
NTNU-run1	0.0724	0.5794	0.23	0.24	0.18	0.1517
Kasetsart-kas16-EXT	0.0585	0.3756	0.1625	0.1313	0.1125	0.1021

**Table 4.** Best performing runs (only showing one run per group) based on MAP over the 20 INEX 2009 and 2010 ad-hoc topics.

Run	MAP	1/rank	P@5	P@10	P@20	P@30
Renmin-LDT2012_adhoc_ruc_comb07	<b>0.3195</b>	<b>0.7655</b>	<b>0.416</b>	<b>0.306</b>	<b>0.231</b>	<b>0.188</b>
Amsterdam-inex12LDT.adhoc.baseline.LM	0.2704	0.7615	0.4	0.294	0.208	0.1673
Otago-ou2012pr09	0.3264	0.741	0.396	0.318	0.233	0.188
NTNU-run1	0.3014	0.7099	0.42	0.316	0.228	0.1733
Kasetsart-kas16-PHR	0.1434	0.7	0.18	0.16	0.085	0.0633
EMSE-run-085	0.3157	0.6979	0.424	0.316	0.235	0.186
MPI-submission	0.1618	0.5991	0.2732	0.1829	0.1061	0.0772

**Table 5.** Best performing runs (only showing one run per group) based on 1/rank over the 50 Jeopardy topics.



**Fig. 2.** Best run by each participating institute measured with MAP

Jeopardy topics, even though some of them submitted the runs to the Jeopardy task not to the Ad-hoc Task. We observe that Renmin University of China (0.7655) runs the first in terms of the mean reciprocal rank ( $1/\text{rank}$ ), but University of Otago (0.741) has the best MAP. The second best scoring team in terms of  $1/\text{rank}$  is University of Amsterdam.

## 7 Conclusions and Future Work

The Linked Data Track, which was a new track in INEX 2012, was organized towards our goal to close the gap between IR-style keyword search and semantic-web-style reasoning techniques. The track thus continues one of the earliest guiding themes of INEX, namely to investigate whether structure may help to improve the results of ah-hoc keyword search. As a core of this effort, we introduced a new document collection, coined *Wikipedia-LOD v1.1*, of XML-ified Wikipedia articles which were additionally annotated with RDF-style resource-property pairs from both DBpedia and YAGO2. This document collection serves as the basis for three tasks: i) the Ad-hoc Retrieval Task, ii) the Faceted Search Task, and iii) a new Jeopardy Task, which were all held as part of this year’s Linked Data Track. We believe that this track encourages further research to-

wards applications that exploit semantic annotations over large text collections and thus facilitates the development of effective retrieval techniques for the same.

## References

1. Media Wiki.  
<http://www.mediawiki.org/wiki/MediaWiki>.
2. SPARQL FullText, W3C Working Draft.  
<http://www.w3.org/2009/sparql/wiki/Feature:FullText>.
3. Wikipedia To XML Extension.  
<http://www.mediawiki.org/wiki/Extension:Wiki2xml>.
4. S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. G. Ives. DBpedia: A Nucleus for a Web of Open Data. In *ISWC/ASWC*, pages 722–735, 2007.
5. J. Hoffart, F. M. Suchanek, K. Berberich, E. Lewis-Kelham, G. de Melo, and G. Weikum. YAGO2: exploring and querying world knowledge in time, space, context, and many languages. In *WWW (Companion Volume)*, pages 229–232, 2011.
6. V. Lopez, V. S. Uren, M. Sabou, and E. Motta. Is Question Answering fit for the Semantic Web?: A survey. *Semantic Web*, 2(2):125–155, 2011.
7. A. Schuth and M. Marx. SPARQL FullText, W3C Working Draft.  
<http://staff.science.uva.nl/~marx/pub/INEX/facetedtaskproposal.pdf>.
8. Q. Wang, G. Ramírez, M. Marx, M. Theobald, and J. Kamps. Overview of the INEX 2010 Data Centric Track. In *INEX*, volume 7424 of *Lecture Notes in Computer Science*, pages 118–137. Springer, Heidelberg, 2011.