



ECP-2008-DILI-528001

EuropeanaConnect

D1.5.1 – Report on tests and evaluation of semantically based functionality

Deliverable number/name	<i>D1.5.1</i>
Dissemination level	<i>Public</i>
Delivery date	<i>26 August 2011</i>
Status	<i>v.1.0</i>
Author(s)	<i>Kirubel Legasion (DNB) Elaheh Momenti Roochi (UW) Antoine Isaac (VUA) Marlies Olensky (HUB) Lars G. Svensson (DNB)</i>



eContentplus

This project is funded under the eContentplus programme, a multiannual Community programme to make digital content in Europe more accessible, usable and exploitable.



europæana
connect

D1.5.1 – Report on tests and evaluation of semantically based functionality



co-funded by the European Union

The project is co-funded by the European Union, through the **eContentplus** programme

<http://ec.europa.eu/econtentplus>



EuropeanaConnect is coordinated by the Austrian National Library

Executive summary

This deliverable is part of the collaboration effort within EuropeanaConnect WP1 partners focusing on testing and evaluating the semantically based functionalities of the Europeana.eu portal application.

The most effective way to evaluate the semantic functionality is to test the quality and competitiveness of the search feature. Hence, two search platforms were considered in this testing task; the Normal Solr and the SKOS Solr.

The Normal Solr means data are collected and organized in the form of Solr index dump whereas in the case of SKOS Solr, these data were augmented with data (mapped) from controlled vocabularies (in this case, the German SWD and the French Rameau).

An evaluation tool has been developed to enable queries to be sent to both SKOS Solr and Normal Solr search engines and list out the different results for evaluation. Initial tests have demonstrated the efficacy and stability of the tool and the proposed methodology.

Quantitative and qualitative analysis was carried out on the evaluation results. However, the quantitative analysis doesn't have statistical significance since very few test users have participated in the evaluation process. A relatively higher number of user responses will be required in order to obtain statistically significant results from the quantitative evaluation of the indices.

The quantitative analysis showed that Normal Solr has higher precision of good results than the SKOS Solr index and more noise (results that are evaluated as bad) was observed in SKOS Solr server than in Normal Solr index.

The qualitative analysis showed that some of the results for a given query sent to SKOS Solr were related to the query not directly but semantically.

The summary of this test predicted two factors that can play role to improve the semantic functionalities drastically:

- a) Integrating more controlled vocabularies and thesauri in the SKOS-based enrichment process
- b) Enhancing the cleanness and completeness of metadata at the ingestion or enrichment phase.



Table of Contents

Executive summary	3
1. Introduction.....	5
2. T1.5 Evaluation Tool	6
2.1 General Structure of T1.5 tool.....	6
2.2 Functional Requirements of T1.5 Evaluation tool.....	7
2.3 Features of the tool	7
2.3.1 Admin GUI (to upload queries).....	7
2.3.2 GUI for Evaluators	7
3. The enrichment process of the Solr.....	9
3.1 Three steps in the Enrichment process	9
4. Evaluation Methodologies and Procedure	10
4.1 Methodologies	10
4.2 Evaluation Procedure.....	11
5. Evaluation Results and Analysis.....	12
5.1 Quantitative evaluation results analysis.....	12
5.2 Precision.....	19
5.3 Qualitative evaluation results analysis.....	20
6. Conclusion.....	20
Acronyms.....	21
Annex - I	22
Annex – II	23

1. Introduction

The Semantic Layer developed in EuropeanaConnect provides uniform machine-actionable and web-enabled access to semantically interlinked resources such as controlled vocabulary data collected from Europeana Partners (in the form of thesauri, heading lists, classification schemes and similar other detailed metadata information). These data can be used to make more semantics-intensive functionalities available for Europeana users.

These controlled vocabularies are converted to the RDF format to form the Semantic Layer using the SKOS model. SKOS¹ (Simple Knowledge Organization System) is a formal language designed for representation of thesauri, classification schemes, taxonomies, subject-heading systems, or any other type of structured controlled vocabulary. SKOS is designed as a modular and extensible language, and in a way that its use and implementation should be as simple as possible. This allows the semantic layer of Europeana to have a uniform representation of the concepts, present in the vocabulary and also helps to semantically align them.

Solr² is an open source enterprise search platform from the Apache Lucene project which is a powerful tool for full text searches from a given dump file. The data collected from Europeana.eu partners is organized in the form of Solr dump data.

For the evaluation of the semantic layer, we compared the search results from two Solr instances: the “Normal Solr” instance and the “SKOS Solr” instance. Normal Solr search refers to the act of searching texts (queries) from Europeana’s Solr dump data. SKOS Solr search refers to the act of searching texts (queries) from the same dump data enriched with mappings from SKOSified controlled vocabularies.

We used a small evaluation tool that helps to send queries to both Normal Solr and SKOS Solr servers and compares the first 48 documents in both results and filters out the differences in the search results. It is important to note that the results from both Solr indices are ranked. Therefore, the first results are presumably the most relevant ones. The Europeana portal provides a maximum of 12 results per page for a given search query. Hence, it is assumed that the portal should provide the most relevant results in the first four pages.

Evaluators would then mark these differences as “Good”, “Acceptable”, “Bad” or “Don’t Know” by evaluators with respect to the corresponding searching texts (queries). In addition to that, this tool enables evaluators to write brief description of their observation or suggestion on each specific result.

¹ http://en.wikipedia.org/wiki/Simple_Knowledge_Organization_System

² http://en.wikipedia.org/wiki/Apache_Solr

2. T1.5 Evaluation Tool

This tool was developed by the DNB and is currently available online in the Europeana sandbox:

<http://europeana-t15.isti.cnr.it:8080/eConnect/>

2.1 General Structure of T1.5 tool

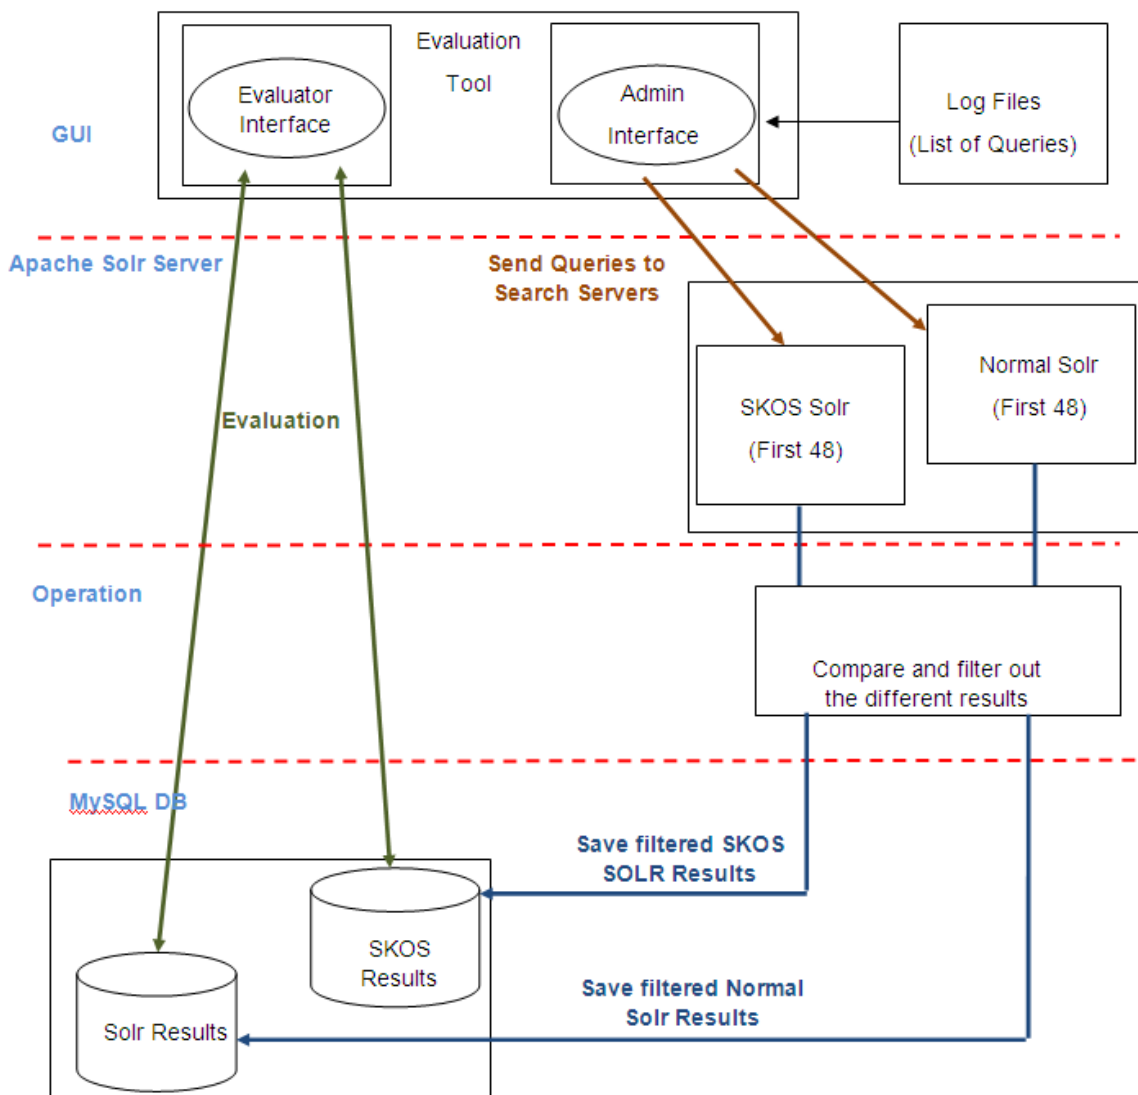


Fig.1: General structure of T1.5 Evaluation Tool

2.2 Functional Requirements of T1.5 Evaluation tool

The tool

1. Should be a web application in order to enable flexible asynchronous technical communication
2. Should be able to send queries to Apache Solr search servers; one with the normal data (Normal Solr) and the other one with enriched data (SKOS Solr)
3. Should have a GUI to upload queries, execute and store the results in local database
4. Should be able to filter out and remove dead links (i.e., the Europeana_URIs)
5. Should be able to compare the first 48 documents from both results based on the unique identifier of each document (i.e., the Europeana_URI)
6. Should be able to filter out and list the different documents from both sides on another GUI
7. Should provide a platform for evaluators to evaluate results. For the sake of simplicity, there are only 4 rating levels, as “Good”, “Acceptable”, “Bad” or “Don’t Know”
8. Should enable evaluators to write brief notes for each result
9. Should enable evaluators save and continue evaluation some other time

2.3 Features of the tool

There are two graphic user interfaces: Admin GUI - to upload queries and GUI for evaluators.

2.3.1 Admin GUI (to upload queries)

The GUI to upload queries is available on the following link:

<http://europeana-t15.isti.cnr.it:8080/eConnect/LoadQueries.html>

1. This GUI helps to upload the pre-selected queries (which were stored in txt or log file). Each log file can have ‘N’ number of queries list. There must be one log file for each evaluator.
2. When the log files are uploaded, an ID is also created automatically for each evaluator (that corresponds to the uploaded list of queries), which will enable the evaluator to log in and evaluate the results with respect to the given query.
3. When uploading and executing the queries into both SKOS Solr and Normal Solr is completed, the results will be stored in local server-side database, followed by a confirmation message on the html page.

2.3.2 GUI for Evaluators

The GUI for evaluators is available on the following link:

<http://europeana-t15.isti.cnr.it:8080/eConnect/>

1. The evaluator can login with the ID created on the admin/uploading interface
2. The evaluator will have access only to the results of the queries that were listed in his/her the log file.
3. The results show not only the link (to redirect the evaluator to the object) but also the Title, Date, Creator, Description, Provider and also the Image.



4. The evaluator can evaluate the resulting document as ('Good', 'Acceptable', 'Bad' and also 'Don't Know'). Moreover, a multi line text area is available to write additional notes.
5. The evaluator has the possibility to save and resume evaluation process next time from that point onward.

3. The enrichment process of the Solr

In order to prepare the enrichment process of the Europeana Solr instance, the following activities have been carried out at the University of Vienna (UW):

- For testing different languages, two SKOS files were selected; Rameau³ (in French) and SWD⁴ (in German), which were provided in the Semantic Layer.
- The triple syntax of these SKOS files was ingested into a Virtuoso⁵ triple store, which is available at UW.

3.1 Three steps in the Enrichment process

1. Find the highest score matches between the concepts in the SKOS files and Europeana documents in the Solr instance and extracting all related labels for each Europeana document from the related concepts.
2. Add the additional labels for each Europeana document (which were retrieved in the first step) using the mapping between SKOS files (we used a mapping file between SWD and Rameau).
3. Update the documents in the Solr index using the extracted labels. In order to keep track of the search results in future (to understand, which label from which SKOS file assists the search), we updated the Solr instance with two new enrichment fields, one for Rameau related labels and one for SWD related labels.

For all sub-processes and for temporarily keeping the results, a Virtuoso triple store has been used.

Result of the enrichment process:

- In the first step, by limiting the search results to the 10 highest score matches, 87.775 matches were found between the SKOS concepts and Europeana objects (in the Solr index).
- In the second step, by performing the mapping between the concepts, 25.675 matches were added. This means, a total of 113,450 matches were found.
- In the third step, for each match all the related labels were extracted and added as an enrichment fields into the Solr index and each document in the Solr index was updated.

Some of the documents in the Solr index were enriched with related labels in different languages. The first and second step ran quite fast but the third step required relatively more time in order to update all matches.

³ <http://sandbox08.isti.cnr.it/econnwp1svn/econnectwp1/trunk/vocs/BnF/>

⁴ <http://sandbox08.isti.cnr.it/econnwp1svn/econnectwp1/trunk/vocs/SWD/>

⁵ <http://virtuoso.openlinksw.com/dataspace/dav/wiki/Main/>

4. Evaluation Methodologies and Procedure

The methods applied to analyze the evaluation results from these two search servers were very simple and clear. Moreover, simple guiding procedures were defined (as shown below in section 4.2) in order to further clarify the evaluation process.

4.1 Methodologies

Two general approaches/methods to analyze the evaluation results:

- 1) **Quantitative analysis:** There are again two ways to apply quantitative analysis under this condition;
 - I. The first way is rather basic. Simply observe which index returns more results.
 - II. The second way is to have each result evaluated by the evaluators as “good”, “acceptable”, “bad” or “don’t know”. Then, by taking the percentage of the evaluation results to the sum of all results in that search index, one can have a better image of the index’s performance. In other words, the more the results are evaluated as “good”, the better the search performance is.
- 2) **Qualitative analysis:** This is a non-trivial analysis. The concept behind this method is to select some queries and evaluate all the resulting documents if it is possible to find semantic relations. For instance, if the given query is “airplane” and if some of the resulting documents have titles as “Flugzeug” in German or “avion” in French or “flygplan” as in Swedish, then it is possible to conclude that the SKOS Solr search is performing better than the Normal Solr search.

4.2 Evaluation Procedure

1. Input queries were collected from two sources, provided by real users. These sources are:
 - a. Europeana portal queries
These are queries that came from Europeana portal, as a result of Log file analysis done by DJ Clark in WP3. That includes both search box queries (typed by the user) and pre-defined queries that came from the Europeana's carousel or "people are thinking about" items. Queries using Lucene's field query syntax, including the advanced search queries (resulting in string like "text:" being appended in the query) were removed.
 - b. Google queries
These are queries that came from Google, which were gathered and cleaned by Jan Molendijk from the Europeana Office. That set was several months old, but is still considered as relevant.

Both queries sets were combined together as query pool, afterwards it was manually cleaned and duplicates were removed.

2. Five evaluators from DNB and HUB selected 50 queries each (a total of 250 queries) from the combined query pool. Please refer to Annex-I to see the 250 queries selected by the evaluators.
3. These selected queries were uploaded to the evaluation tool using the Admin/Load Queries GUI.
4. After the queries were uploaded, the evaluation tool sent out each query from the list to both Normal Solar and SKOS Solar indices and stored the first 48 corresponding results in buffer memory. The reason why only the first 48 results were taken into consideration is that the Europeana portal application displays a maximum of 12 results per page for a given request and one factor to determine the competitiveness and quality of Europeana's search engine is to find most relevant results in the first four pages, i.e., a maximum of 48 results.
5. The evaluation tool compared the results that were stored in buffer memory and filtered out only the differences.
6. The filtered out different results were stored in a local database and linked to the profile of the evaluator who has selected these queries initially.
7. The evaluators can then login to their profile and proceed with the evaluation. This evaluation process is a subjective process where each evaluator takes into account his/her background knowledge, experience and also help of additional supportive references, such as Internet resources, to determine whether that specific result would be a good result or not.

5. Evaluation Results and Analysis

5.1 Quantitative evaluation results analysis

There were a total of 45 queries out of 250 that had different results in Normal Solr and SKOS Solr configurations, as shown in Table 1.

Table 1: Queries with different results

	Query	SKOS Solr Result	Normal Solr Result
1	cuba	✓	✓
2	jules verne	✓	✓
3	Da Vinci	✓	✓
4	art deco	✓	✓
5	costume oriental	✓	✓
6	film	✓	✓
7	films	✓	✓
8	robe	✓	✓
9	astronomy	✓	✓
10	claudes monet	✓	✓
11	malta	✓	
12	tahiti	✓	✓
13	ukraine	✓	✓
14	charles darwin	✓	✓
15	Chili	✓	✓
16	dali	✓	
17	museum	✓	
18	roma	✓	✓
19	griechische plastik	✓	
20	Finland	✓	✓
21	willy brandt	✓	
22	tibet	✓	✓
23	typographie	✓	✓
24	Scotland	✓	✓
25	napoleon	✓	✓
26	Anglican church	✓	✓
27	aphrodite	✓	✓
28	art nouveau	✓	✓
29	Cambridge	✓	✓
30	halsring	✓	✓
31	kafka	✓	✓
32	Love OR Liebe OR Amour OR amore	✓	✓
33	salvador dali	✓	
34	alchemie	✓	✓
35	arab	✓	✓
36	bible	✓	✓
37	book	✓	✓
38	voltaire	✓	✓
39	archaeology	✓	✓
40	athena		✓
41	chopin	✓	
42	costume renaissance	✓	✓
43	Daguerreotype	✓	✓
44	david goliath		✓
45	Louis XIV	✓	

This tool sends queries to both Normal Solr and SKOS Solr and takes the first 48 results for comparison. After comparing these 48 results from both sides, it filters out only the differences. After filtering out the differences, the tool checks the Europeana_URIs (links) of these differences whether they would redirect the evaluators to the actual objects or not. If the links are dead (i.e., they wouldn't redirect evaluators to the objects), then the tool removes these results away in order to reduce unnecessary efforts that would be spent by evaluators.

As shown in Table 1 above, there are 8 queries (malta, dali, museum, griechische plastic, willy brandt, salvador dali, chopin and Louis XIV) that have results in the SKOS Solar search but not in the Normal Solr search.

On the other hand, there are only two queries (Athena and david goliath) that have results in the Normal Solr search but not in the SKOS Solr search. Theoretically, it is expected that the results in Normal Solr should also be observed in SKOS Solr. But as explained above, such occurrences can be observed if the SKOS Solr search results have dead links.

According to our findings SKOS Solr search delivers better results than Normal Solr search.

Table 2 shows the exact evaluation results of each document for all queries, in both SKOS Solr and Normal Solr.

Table 2: Evaluation Results

ID	Query	Total SKOS-Solr Results	Total Normal-Solr Results	SKOS Solr Evaluation				Normal Solr Evaluation			
				Good	Accept.	Bad	Don't Know	Good	Accept.	Bad	Don't Know
1	cuba	1	4	1				4			
2	jules verne	1	1	1					1		
3	Da Vinci	1	1			1					1
4	art deco	1	1			1				1	
5	costume oriental	1	1		1				1		
6	film	1	1	1					1		
7	films	1	1	1					1		
8	robe	1	1		1				1		
9	astronomy	9	5	8		1		5			
10	claudes monet	1	1	1				1			
11	malta	1	0	1							
12	tahiti	2	2	2				2			
13	ukraine	17	15	17				15			

14	charles darwin	1	1	1				1			
15	Chili	16	20	16				20			
16	dali	6	0			6					
17	museum	2	0	1							
18	roma	1	1	1				1			
19	griechische plastik	1	0		1						
20	Finland	9	3		9				3		
21	willy brandt	1	0	1							
22	tibet	1	1	1				1			
23	typographie	2	1	2				1			
24	Scotland	3	3	3				3			
25	napoleon	1	1			1		1			
26	Anglican church	1	1	1					1		
27	Aphrodite	1	1	1				1			
28	art nouveau	10	7	6	4			7			
29	Cambridge	4	4		4				4		
30	halsring	4	4	4				4			
31	kafka	1	1	1					1		

As

32	Love OR Liebe OR Amour OR amore	2	9	2				4	5		
33	salvador dali	1	0	1							
34	alchemie	4	4	4				4			
35	arab	13	17		8	5		4	12	1	
36	bible	15	10	13	1	1		6	2	1	1
37	book	9	9		9				8		1
38	voltaire	6	6	6				6			
39	archaeology	4	2		4				2		
40	athena	0	1					1			
41	chopin	1	0	1							
42	costume renaissance	1	1			3					1
43	Daguerreotype	1	0	1							
44	david goliath	0	9					9			
45	Louis XIV	5	5	4	1			2	2	1	
Total		165	156	105	43	17		103	45	4	4

As indicated in Table 2 above, the total number of documents specific to the SKOS Solr for the 45 queries is 165 whereas Normal Solr provided 156 documents that are not provided by the SKOS Solr.

The reason why Total SKOS Solr results are different from the Total Normal Solr results, as shown in Table 2, is because the tool filters out dead links before listing them.

For converting the above values into percentage see table here below.

Table 3: Summary of evaluation results

	Good	Accept.	Bad	Don't Know
SKOS Solr results (out of 165)	63.7%	26%	10.3%	
Normal Solr results (out of 156)	66%	28.8%	2.6%	2.6%

The above numeric figures in Table3 don't have statistical significance since such kind of testing methodologies are very subjective and in this case, only five evaluators have participated in the testing process.




Fig.2: Evaluation results for **SKOS Solr**

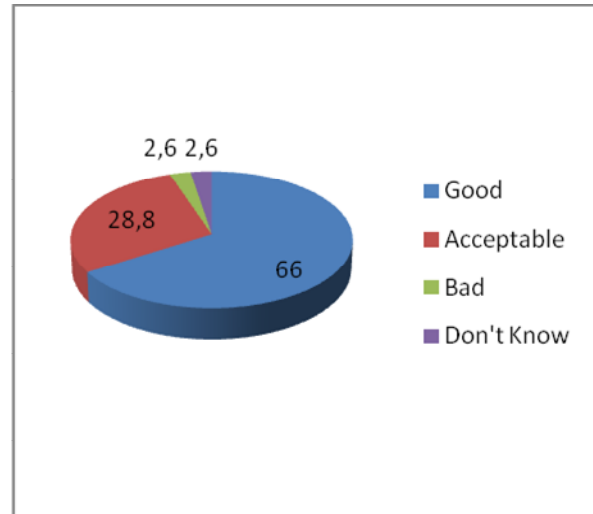


Fig.3: Evaluation results for **Normal Solr**

The above diagrams (Fig.2 and Fig.3) show that the percentage values of documents that are rated as “good” in Normal Solr is greater than those in SKOS Solr. On the contrary, the percentage values of documents rated as “bad” in Normal Solr is less than that of SKOS Solr.

A possible explanation for this discrepancy is observed as follows: For instance, for the query “dali”, SKOS Solr has returned six documents whereas Normal Solr has found nothing. But SKOS Solr has tried to relate “dali” with “daly” and all the six documents contained “daly”. Unfortunately, there was no other metadata information to help the requester (in this case, the evaluator) to see if these two results are somehow semantically related or not. Hence the evaluator rated all the six documents as “Bad”. The same problem can also be observed with the results of other queries such as “arab”. This observation indicates that “Bad” evaluations may have ambiguous implications for enriched indices, and that future experiments should be designed to take this factor in to account.

5.2 Precision

Precision can be seen as a measure of exactness or quality.⁶ Precision can be calculated as the number of relevant documents a search retrieves divided by the total number of documents retrieved.

For example,

Precision of good = $P_{\text{good}} = (\sum (\text{total good results})) / (\sum (\text{total good results}) + (\text{total accept results}) + (\text{total bad results}))$

In the case of SKOS Solr,

The precision of “good”

$$P = 105 / (105+17+43) = \mathbf{63.6\%}$$

The precision of “acceptable”

$$P = 43 / (105+17+43) = \mathbf{26\%}$$

The precision of “bad”

$$P = 17 / (105+17+43) = \mathbf{10.3\%}$$

In the case of Normal Solr,

The precision of “good”

$$P = 103 / (103 + 45 + 4 + 4) = \mathbf{66\%}$$

The precision of “acceptable”

$$P = 45 / (103 + 45 + 4 + 4) = \mathbf{28.8\%}$$

The precision of “bad”

$$P = 4 / (103 + 45 + 4 + 4) = \mathbf{2.5\%}$$

The above results indicate the relative quality of each evaluation results compared to the overall results of the same Solr index.

⁶ http://en.wikipedia.org/wiki/Precision_and_recall

5.3 Qualitative evaluation results analysis

For instance, for the query “book”, SKOS Solr provided results that contained documents with titles as “Schreib-Formularbuch...” and/or “Stempel-bok ...”. These resulting documents are supposedly related to the given query semantically. These results would not have been provided by the Normal Solr search. This indicates that the enhanced Solr index can provide added value to users by including related documents that would not have been found without the semantically interlinked subject vocabularies.

6. Conclusion

The project has produced a tool for the evaluation of enriched indices. Initial tests have demonstrated the efficacy and stability of the tool and the proposed methodology. However, a significantly higher number of user responses will be required in order to obtain statistically significant results from the quantitative evaluation of the indices.

The qualitative analysis showed that some of the results for a given query sent to SKOS Solr were related to the query not directly but semantically.

This evaluation tool filters out only those queries that have different results between SKOS Solr and Normal Solr. That means, only 45 queries out of 250 queries have different results, which would be around 20%. The rest of the results, which are around 80%, are alike. This is because currently, only two vocabularies, namely SWD and Rameau (as indicated in section 3 of this report) are used to SKOSify the dump. Hence, many of the search results are quite similar.

It is our expectation that the following two points would enhance SKOS Solr results:

- Integrating more controlled vocabularies and thesauri in the SKOS-based enrichment process is needed;
- The enrichment process should be enhanced so that lower noise is brought in the results (i.e., documents are not mapped to irrelevant concepts).

This is an area where more research and experiments are needed in order to verify the expectations.

Acronyms

	Abbreviated	Full (unabbreviated)
1	DNB	Deutsche Nationalbibliothek (German National Library)
2	HUB	Humboldt-Universität Berlin
3	UW	Universität-Wien

Annex - I

Evaluators' information and their institutes

	Evaluator's Full Name	Institute
1	Anke Meyer	DNB
2	Stephanie Glagla-Dietz	DNB
3	Helga Karg	DNB
4	Marlies Olensky	HUB
5	Steffen Hennicke	HUB

Annex – II

The 250 queries selected by the 5 evaluators are as follows:

Coco Chanel	astronomy	mubarak badenweiler	1.weltkrieg	belagerungsgerät
cuba	claudes monet	meissen oktober	alicante	china
jules verne	dresden	mondsichelmadonna	Cezanne	dürer
kostüm papageno	faschingsdekoration	ortelius	festungsplan	eisschnelllaufanzüge
leipzig bahnhof	gewürzverpackung	pionierpalast dresden	frühlingslandschaft	japan
Puerto Rico	handpflug	paravent	kopenhagen	klimt
quijote	johann sebastian bach	robotron fernseher	sanssouci	london
Da Vinci	malta	schuhpresse	taschenuhr	marseille
dante	Maria Montessori	griechische plastik	van gogh	nietzsche
DDR Plakate	mikroskop	gobelintechnik	warschau	Posen
der lenz ist angekommen	prague	fischeneisen	napoleon	Rugby
ein engel kommt nach				
babylon	schloss güterfelde	Finland	kinderpostamt	torgau
			Gustav II. Adolf, König von	
mode 1951	tahiti	brigitte lahaie	Schweden	tripolis
bauhaus	ukraine	armorial	gallionsfigur	worms
granada	zeitung 1935	bach	corfu	Zweiter Weltkrieg
5 tage 5 nächte film	blockhaus	steinsalzgewinnung	Adressbuch	alchemie
amanda lear	böhmen	soldatenfibel	Anglican church	arab
anne marie peysson	borsten	typografie	aphrodite	augsburg
arles	Breslau	unfallwagen	aristoteles	bible
art deco	charles darwin	vauxhall walk	art nouveau	book
barcelona	Chili	Völkerschlacht	Benjamin Franklin	brazil
brest france	cyprus	walzenständer	british museum	burgen
carmen cartellieri	dali	zeitung	Cambridge	console
cordoba	Danzig	willy brandt	cemetery	serbia
costume oriental	deutsches heimatwerk	zahnradkonstruktion	Christian Dior	siegel
Dior	eckbank	Wappenbuch	cleopatra	sonet
fashion	famagusta ghost town	wandvertäfelung	faschingskostüme dresden	timor
film	fasching 1958	waggonbau görlitz	frau ohne kuss	video
films	fastnachtsumzug freiburg	verviers	gummistiefel	violin

Galicia	gartenstadt hellerau	travestie	halsring	voltaire
garnhaspel	Germanische Nationalmuseum	toledo	hercules	13th century
guitar	gewürzverpackungen	toilettentisch	jerusalem	action man
jugendstil	greifswald	thonet	kafka	alt breslau
leggings	grünberger wochenblatt	tibet	korsett	archaeology
Madrid	hessisches landesvermessungsamt	wool winder	laufbahn der sonne	athena
modes OR fashion	highland home industries	Zeus	leonardo da vinci konstruktionen	bajonett
nordmark film	hochschwarzwaldstraße	weckglasöffner	Love OR Liebe OR Amour OR amore	beatles
ornament	hundeschlittenrennen todtmoos	uniformstiefel	malaga	berchtesgaden
papageno kostüm	James Bond	typographie	mary stuart schiller	brecht
pionierkleidung	jean-jaques rousseau	tiroler anzug	Mozart	caspar david friedrich
reformkleid	kaffeeschütte	turnhalle	new zealand	chopin
robe	karte	taufkleid	olga picasso	Constantinople
schal	königsberg	Scotland	papagena kostüm	costume renaissance
schmuckkamm	lachenhäusle	soap manufacturer	pommersche zeitung 1902	Daguerreotype
			porzellanmanufaktur meißen	
sevilla	lima, peru	siebmacher	oktober	danziger zeitung
stroschuhe freiburg	museum	stereoskop	poseidon	david goliath
	naturräumliche gliederung			
the rolling stones	deutschlands	siegesdenkmal leipzig	rittertugenden	DaVinci
tradition design	roma	stuck	salvador dali	Louis XIV
udo lindenber	sächsische neueste nachrichten	ross	spreeverlauf	lucas cranach
valencia	Schlesien	ring	stolp landesarchiv	Ludwig von Bayern