LISA*: Towards a multi-property comparison interface for linked cultural heritage sources

Alia Amin[1], Michiel Hildebrand[1], Jacco van Ossenbruggen[1,3], Annelies van Niespen[4], Lynda Hardman[1,2]

[1]Centrum Wiskunde & Informatica (CWI), Amsterdam, NL,
[2]University of Amsterdam (UvA), NL,
[3]VU University, Amsterdam, NL,
[4]Digital Heritage Netherlands (DEN), Den Haag, NL.

Abstract. Comparison search is an information seeking task where users need to examine similarities and differences among individual or groups of objects. It is currently only supported for comparing individual objects in homogeneous data sets. In this paper, we discuss comparison search in cultural heritage, a domain characterized by large, rich and heterogeneous data sets, where different organizations deploy different schemata and terminology to describe their artifacts. These differences make meaningful comparisons across different museum collections much harder. To support comparison search in heterogeneous data, we designed a prototype based on use cases identified in previous work. The prototype allows the users to compare sets of heritage artifacts using different data visualizations, deploying different strategies to deal with the heterogeneity of the underlying data. We used this prototype to conduct a preliminary evaluation on the supported features with cultural heritage experts. Based on our evaluation, we identified issues that are important to our users, such as trust in the data alignment and query formulation improvement for set selection.

Key words: Comparison search, sensemaking, cultural heritage

1 Introduction

In an in-depth study of information seeking needs of cultural heritage experts, we found that Information Gathering tasks are a significant part of their daily work [2]. Of the different types of Information Gathering, Comparison search is one of the most challenging tasks. Comparison search in an information seeking task that involves comparing differences and similarities between objects or sets of objects [2]. It requires a complex analysis from different views to understand the characteristics of multiple-collections. Unfortunately, there is little explicit support for comparison search on linked collections of annotated images, and in particular the cultural heritage collections on which our work is based.

* Linked data Interface for Sensemaking and Analysis

Consider an art historian doing research on the rich diversity of Dutch painting collections. These collections are stored in different cultural heritage institutions, e.g. state owned, or independent organizations. To save time, she might use a search aggregator to help her search across their online archives. But the work has only just begun, as the search results provide a list of many different paintings corresponding to the search queries. The historian needs to investigate distinguishing properties of these paintings (e.g. artists, materials used, artstyles) in order to fully understand the characteristics of each collection. Unfortunately, different institutions use different vocabularies to describe their collections. Therefore, she must first be familiar with the different vocabularies to be able to compare, interpret and match the vocabularies belonging to the different collections. The next challenge is to organize, assemble and analyze a vast amount of multidimensional information from her search results coherently in away that she can derive meaningful and accurate conclusion.

Related to the scenario above, we highlight the main challenges users face. First, the specific level of expertise required to understand the vocabularies belonging to different cultural heritage institutions. Different museums are likely to have their own in-house thesauri in addition to general ones. Learning the vocabularies of different museums requires time and assistance from experts from each museum. Second, the time and effort required to gather pieces of information. Most museum archive and digital libraries provide only two features: searching for and browsing objects within a collection. These interfaces are insufficient for finding answers to complex information tasks, such as comparison search. Consequently, pieces of information have to be gathered and organized manually, through laborious repetitive work. Such repetitive work could have been computed in a fast and cheap way automatically. Moreover, with the appropriate interface and visualization, it should be possible to discover new ways to look at inter-connected collections and assist complex comparison between them. Thus, allowing the expert to carry out the in-depth analysis work faster.

In this paper we describe our approach to support comparison search between different heterogeneous cultural heritage sources, the challenges that we face in our approach and the results of a first evaluation with cultural heritage experts. The contribution of this paper is to establish the design requirements for the information visualization and interaction to support comparison search across semantically linked cultural heritage sources. One of our motivations is to understand how recent developments in the Semantic Web can support comparison search in domains where there are rich heterogeneous structured data, such as cultural heritage.

The paper is structured as follows. We first describe related work. Next, we discuss functional requirements of a comparison search. Afterwards, we describe our comparison search application, including the technology infrastructure, data sets and interface. We focus the discussion on the supported features at the interface. Afterwards, we discuss our preliminary evaluation with cultural heritage experts. We end with a discussion about our future work for the interface and a conclusion.

2 Related Work

To support comparison across heterogeneous cultural heritage sources, we need to select the appropriate visualization and interface features, and to identify the functional requirements for such task.

Information visualization — Visualizations should not only present information but also invite users to actively engage in the direct manipulation of parameters, and allow users to discover new knowledge, find patterns, observe changes and identify characteristics [19]. There are a number of studies that used information visualization in digital libraries, catalogs and information aggregation systems. These studies used different types of visualizations, such as graphs e.g. [9, 16], hyperbolic trees e.g. [10], maps e.g. [1], tables e.g. [14, 17], and scatter plots e.g. [10, 11]. A graph is one type of visualization used in several aggregator systems, e.g. WebMaster [16], MUSA [9]. In these systems, the graph nodes usually represent entities (e.g documents, events, or concepts) and the graph edges represents the relationship between two entities. A similar visualization is the hyperbolic tree but with a distorted view e.g. in [10]. Graphs and hyperbolic trees are typically used to present relationships between information or when there is a need to navigate between entities.

The ability to display multiple properties is an important criteria in choosing a visualization for the comparison search interface [15, 8]. Some visualizations can only show the values of one property e.g. map, or bar chart, while others are able to display more simultaneously, such as scatter plots (2 properties) and tables (multiple properties). A map visualization can be used to present location-based information. It is widely used online due to the availability of online APIs such as Google Maps[1] and Yahoo Maps[2]. In [1], a map is used to represent aggregated Flickr[3] tags corresponding to physical places/locations. Similar to a map, a bar chart can only show the values of one property. The difference is, however, a bar chart can show values from any given property (generic category). A table has been used in comparison search applications for many years e.g. [4, 15]. The downside of this visualization is that it is only appropriate when there is a single value for each property. A scatter plot is mostly used when there are two categories to display at once, such as in Gapminder and in [10]. The Gapminder application have also demonstrated additional techniques on how to add more information in the visualization by using distinctive physical traits such as size, color, shape or symbol. This technique is useful when there is limited space to visualize different types of information. In Table 1, we summarize the characteristics of different visualizations. In the following paragraphs, we discuss the types of visualization traditionally used in comparison search interfaces.

Comparison search applications— In the cultural heritage domain, there are already applications that support searching across multiple cultural heritage collections, such as the Collectiewijzer[4], ECulture MultimediaN [12], Culture-

[1] `http://maps.google.com`
[2] `http://maps.yahoo.com`
[3] `http://www.flickr.com`
[4] `http://www.collectiewijzer.nl`

Table 1. Characteristics of different visualizations

Requirement	Graph	Map	Bar chart	Table	Scatter plot
Show comparison	no	yes	yes	yes	yes
Show property	multiple	1	1	multiple	2
Show many values for a property	yes	yes	yes	no	yes

Sampo [5]. To date, these systems primarily support three information seeking tasks, namely searching, browsing and exploration. This is also true for most digital libraries systems [13]. Even though important, these tasks are only a few from a wide range of information seeking tasks cultural heritage experts do for their professional work [2]. Another important task is comparison search. Many research on comparison search can be found in the commerce domain, nevertheless, comparison search is an information seeking need that extends in any domain where there is a need to examine the difference between two or more sets [2]. There is a difference between comparison search in the commerce domain and in the cultural heritage domain. In the commerce domain people are typically interested in comparing single objects, whereas in the cultural heritage domain, comparison can be between single objects (e.g. compare two Van Gogh's painting) and also comparison between sets of objects (e.g. compare the paintings in the Rijksmuseum Amsterdam with the paintings in RKD[5] Den Haag). In the commerce domain, optimizing interfaces to support comparison search has long been an important research topic, e.g. in [4, 7, 8, 15].

Several research on product comparison interfaces have highlighted the importance of allowing users to extensively search and browse objects [4, 15]. A study on electronic catalogs [15] emphasizes the importance of providing features that enable incremental object selection. The VOPC interface [8] is designed to enable users to interactively filter objects based on the available properties. These properties are visualized next to each other. Thus, providing the user with an overview of all possible options. The effects of different types of interfaces on user performance have been studied, e.g. in [4, 7]. The experiment in [7] have shown that a table-like interface, e.g. Eureka, would help users solve problems faster, while a scatter plot interface, e.g. Spotfire, is better at helping users find correct answers. Callahan et al. [4] compared an interactive table (InfoZoom) with a hierarchical table interface. They found InfoZoom helps people compare object properties faster than the hierarchical table interface. The hierarchical table interface, however was more pleasant to use. These studies suggest that different visualizations have their own strengths and weaknesses.

Previous research suggests that the user's performance while using a comparison search interface depends on the type of task, the context and the ability of the user to translate the given problem while working with the system [4]. Thus, different domains might have slightly different comparison search requirements. We found no prior work on comparison search in the cultural heritage domain.

[5] Netherlands Institute for Art History, `http://website.rkd.nl/Databases`

2.1 Functional requirements

To our knowledge, there is no application that (a) supports comparison search in the cultural heritage domain, (b) is dedicated for comparison search for sets of objects as opposed to comparing single objects, (c) uses semantic relationships to support comparison search across multi-sources, and (d) enables user to switch between the available properties while comparing sets of objects. Prior work can thus only provide us with rough guidelines for the comparison search interface we wish to provide. Here, we summarize several basic functional requirements distilled from the above literature for a comparison search interface that can support users to search, select and compare (sets of) objects effectively.

First, a comparison search interface should enable users to iteratively search across museum collections and view search results before selecting objects of interest. The search interface should present users with all available facets of the different collections to give an idea about the collections and how to navigate through them. Second, the interface should provide a means for the user to make a selection of specific (sets of) objects to compare. The selection interaction should be simple and there should be feedback from the system to assure the user that a selection has been made. Additionally, it should be possible to undo any action should the user wish to do so. Third, while comparing the selected objects, the visualization should present at one or more properties and enable users to see the all values for the given property. One way to meet all requirements and enable different views is by enabling multiple visualizations. Each visualization technique has its own strengths and weaknesses. It might be useful to use several visualizations that complement each other. Thus, users will benefit from being able to optimally explore different aspects of the information (see Table 1). In the following section, we discuss how we translate the comparison search application requirements in our implementations. Furthermore, we illustrate how to use and interact with the visualizations for comparison tasks.

3 The LISA.ECulture comparison search application

The LISA application[6] is part of a suite of tools developed within the ECulture MultimediaN project [7]. In the following sections we discuss the technology infrastructure, the datasets and the user interface. We focus our discussion primarily on the interface and interaction.

3.1 Technology Infrastructure and dataset

Infrastructure — The LISA application is developed on top of *ClioPatria*, the web infrastructure for search and annotation across heterogeneous collections.

[6] The LISA prototype in progress is accessible at: `http://e-culture.multimedian.nl/lisa/session/compsearch`.

[7] `http://e-culture.multimedian.nl/`

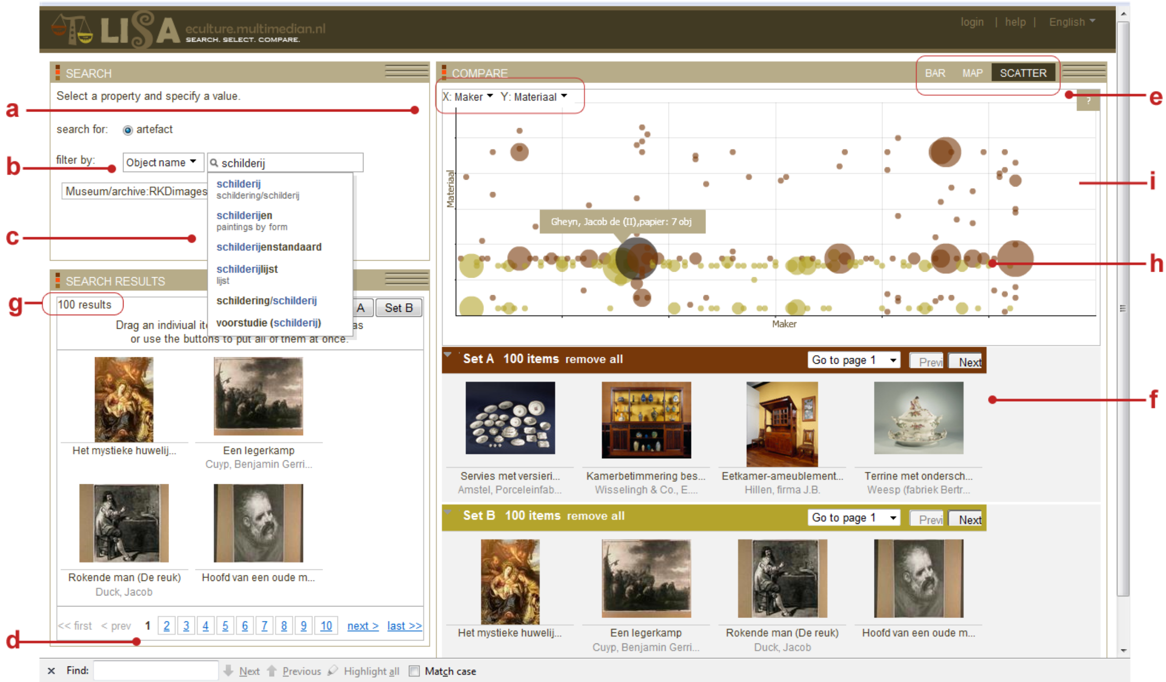

Fig. 1. The LISA interface features: (a) 2D property pull down menu, (b) search filter, (c) auto-completion, (d) search result navigation and browser, (e) visualization options, (f) selection sets, (g) number of search results, (h) indication tooltip, the size of the bubble corresponds to the amount of objects, (i) Scatter plot visualization.

Fig. 2. The bar chart (top) and map (bottom). In the bar chart, objects from Set A and Set B are alphabetically organized according to the creator (Dutch: Maker) in the X-axis, the Y-axis indicates the amount of objects. In the map, selected objects from Set A and Set B are shown according to (location) subject depicted (Dutch: Onderwerp). The balloons indicate the amount of objects and value. Alternating between different properties is done via the pull down menu (a) and (b).

For detailed information on the web server infrastructure and the search strategies across heterogeneous collections, see [12, 21]. Communication between the client and the server is done via a request through the system's APIs. Information is sent back from the server in a form of a JSON object. The implementation of the interface uses (X)HTML, CSS, Javascript and Flash. It is tested on Firefox 3.0.7 browser. The client side visualization widget uses an extension of the Yahoo User Interface widget (YUI v. 2.7.0) and amChart v. 1.6.5 [8].

Dataset — To enable comparison search with LISA, the server needs to host common thesauri, namely IconClass[9], and the Getty vocabularies[10], i.e. the Art & Architecture Thesaurus (AAT), Union List of Artist Names (ULAN) and the Thesaurus of Geographical Names (TGN), as well as collection specific thesauri, such as thesauri from Rijksmuseum and RKD. Table 2 shows the size of collections and thesauri currently used by the application. Collections and thesauri data were converted to an RDF/OWL representation. In principle, to allow information access across collections, specific thesauri are aligned with the common ones. For example, the artists' names in the Rijksmuseum thesauri are linked with artists' names in ULAN. Materials concepts in the Rijksmuseum thesauri are linked with concepts in AAT. Detailed information on the conversion and alignment method of cultural heritage sources in the ECulture MultimediaN project can be found in [20, 18].

Table 2. Thesauri and collections used in LISA[*)]

	source	**size**
Collection:	Rijksmuseum	31.699 objects
	RKD Archive	82.781 objects
Thesaurus:	Rijksmuseum thesaurus	53.525 terms
	RKD thesaurus	11.995 concepts
	TGN	89.000 concepts
	ULAN	13.000 people
	AAT	31.000 concepts
	IconClass	24.331 concepts

[*)] `http://eculture.cs.vu.nl/europeana/www/datacloud.html`

3.2 Interface

Overview — The LISA interface consists of four panels/areas. A search panel is a place where the user formulates queries progressively, and the results are shown in the Search Result panel. In the Selection panel users specify two sets,

[8] `http://www.amChart.com`

[9] `http://www.iconclass.nl`

[10] `http://www.getty.edu/research/conducting_research/vocabularies/`

i.e. Set A or Set B[11]. Visualization panel provides three options: bar chart, map and scatter plot visualization (see Fig.1 and Fig.2). As an example of how the interface works, we compare the art objects from RKD Images archive[12] and the art objects from the Rijksmuseum Amsterdam[13].

Search — The search interface is specially designed to allow users to define a clear set of selection of art objects. Fig. 1 shows how a search is conducted. The user can search simply by typing any keyword and pressing enter (simple search). Alternatively, the search interface is equipped with two features: autocompletion and a property pull-down menu selection (see Fig. 1b). The thesaurus-based autocompletion[14] can be used to guide the query formulation process by selecting one of the autocompletion suggestions. The suggestions are terms coming from different cultural heritage thesauri used by the application (see Table 2). We use an alphabetical ordering for our autocompletion as we found it to be the most effective ordering for loosely structured thesauri [3]. When a property is selected from the pull-down menu, the search will be restricted for the selected property, such as the creator, the material of the art object, or the date of creation. Thus, it is possible to formulate a filtered search on a specific `property:value` pair (advanced search). For example, search on art objects that are paintings (`object name:Schilderij`). The simple search is suitable for collection exploration while the 'advanced search' is used when the user already has specific objects in mind.

Selection — After the user is satisfied with the search results, she needs to select the (set of) art objects to compare. The selection can be done via dragging and dropping an object thumbnail from the search result panel to the selection panel (see Fig.1f). Alternatively, a bulk selection of all search results can be made by clicking the *Set A* or *Set B* button. The search and selection process are typically done iteratively. For example, first the user found results by searching of all paintings from RKD Image Archive (`museum/archive:RKD`), place them on the SET A selection area. Afterwards she made a second search of all paintings from the Rijksmuseum dataset (`museum/archive:Rijksmuseum`) and placed the results in the SET B selection area. At this point, the user already selected two sets of objects to compare.

Visualization — Whenever an object is placed in the Selection panel (either Set A or Set B), the visualization area is always updated. Three types of visualization are available: the map, the bar chart and the scatter plot (see Fig. 1i and Fig. 2). In all visualization, we use color codes to indicate which sets do the elements belong to (either Set A or Set B). These visualizations are chosen because of their ability to present at least one property and multiple values simulataneously (see Table 1). The bar chart visualization can show comparisons

[11] At the time being only two sets comparison is supported. This is extendable as the current application design takes into account the future addition of more sets, however, since computations are carried out client-side, for more than 1,000 objects the interface becomes slow.

[12] `http://english.rkd.nl/Databases/RKDimages`

[13] `http://www.rijksmuseum.nl/`

[14] For information on the thesaurus-based autocompletion design see [6].

between values of *one* property at a time. Fig. 2(top) shows a representation of art objects from both sets. The X-axis represents the objects organized in the selected property:*creator* (Dutch: Maker) in alphabetical order from A to Z. The property pull-down menu shows all available properties for which the objects can be organized, e.g. by dimension height, by date, by creation material, or by depicted subject. The Map visualization is also a one dimensional view. The user can make a selection to view the sets by different location specific properties, such as collection location, location subject, and museum/archive location (see Fig 2b). The values of these properties are typically place name, address, city or country name. The Scatter plot visualization shows the values of two different properties at the same time (see e.g. Fig. 1a.). Fig. 1i shows how sets are distributed by their creator and their material. It is easy to observe that while both sets the creator names are distributed, the material characteristics of both sets, however, are very different. Set A (the RKD collection) tends to have more homogeneous type of material with little variations while Set B (the Rijksmuseum collection) have more variety. In all three visualizations, it is possible to explore the information space by two means (see Fig. 2): a) alternating between multi-property views using the property pull down menu b) detail exploration by zooming, panning and scrolling. With the different visualizations, the user can flip through all distinguishing properties to examine multiple collections simultaneously enabling her to gather quick insights and overview on vast collections that is otherwise extremely difficult to do with a simple search and browse interface.

4 User Interface Evaluation

In this section, we discuss the outcome of our user interface evaluation. To get an early feedback on our application, we conducted a preliminary evaluation in our first design cycle iteration. We went to the cultural heritage experts and demonstrated what the application can potentially do. Our main goals were to gather initial impressions and comments to improve our system and to collect interesting use cases from various experts of how to use our system.

4.1 Procedure

We carried out semi-structured interviews that took place at the participants' working environment. The interview consisted of several parts. We began with an introduction and general demographic questions. Afterwards, we asked questions about participants' activities related to accessing heterogeneous cultural heritage sources and how to compare these information sources. Next, we introduced LISA and gave examples of how the application works. This was followed by a creative brainstorm exercise where the participants were encouraged to generate similar use cases inspired by their own experiences. Finally, the participants were given the opportunity to ask any questions. All interviews were voice recorded for documentation and analysis. In total, seven cultural heritage experts from 3 different institutions took part in the interviews. The participants' average age was 39 years old. They have diverse roles: two senior researchers, 3 curators, 1 art historian, 1 consultant. Most participants have senior positions and have a good overview of the different expert roles and their responsibilities.

5 Discussion

Our preliminary evaluation gave us useful feedbacks. Our cultural heritage experts showed us many new usage of the tools which we had not predicted earlier. This certainly gave us insights on how to improve our prototype. All interview recordings were qualitatively analyzed. In this section, we highlight some of the participants' comments and concerns.

Using the comparison search tool — In our brainstorming session, our participants came up with interesting use cases of how the tool can be useful for their daily work. Here we give three examples. First, comparison search as an educational tool. One expert is also a university lecturer and have an interest in using LISA as a tool to assist students learn about different museum collections. Second, comparison search as a research tool. Cultural heritage researchers can use comparison search as an analysis tool, primarily to explore and understand unfamiliar collections. Third, comparison search as an assessment tool. Curator managers can use this tool to examine the characteristics of different museums. The tool is useful to plan for expert training courses, to conduct risk management and assessment, and to seek expertise from other museums with similar profiles.

Trust in the property and value alignments — LISA is made possible by alignments between the properties and values of the different thesauri. By doing so, LISA presents a new interpretation of information or a new 'derived' knowledge. Since this new knowledge is not provided by the original cultural heritage institutions, this raised concerns about the inner workings of the data conversion and overall quality of the information. One expert asked: *"How do you know if you are matching the museum vocabularies to the right concepts? For example what will happen if this object is made out of (and annotated by) a specific Bamboo which is not in the common vocabulary? How do you deal with this?"*. It is important for these experts to be assured that the quality of the conversion is agreed upon and confirmed by the museums who are the original information provider.

Improve query formulation for set selection — Searching through heterogeneous collections can be quite challenging specially when the user is not familiar with the thesauri used in many different collection. One expert commented that she always receive request from colleagues who wanted to search but could not find the correct term: *"Sometimes I do the searches for other people because they are not familiar with our thesauri, for example if you want to find an object related to the Islamic culture you have to know several terms such as 'Islam' (the religion), 'moslem' (person singular), 'muslim' (person singular, male), 'muslimah' (person singular, female), 'muslims' (person plural)"*. It is important to support and assist users while formulating queries. One way to help users cope with this problem is by providing suggestions on semantically related concepts, for example in the autocompletion interface.

Related to this problem is multilingualism. For example, while searching on the keyword 'Painting', the system should retrieve results of objects annotated by 'Schilderij'. Different collections might be annotated in different languages. The interface should assist users to overcome this language barrier by providing suggestions of alternative spellings, synonyms and alternative language.

Follow up work — Our main focus in the development process is on improving the quality of the data alignments. In order to enable comparison, all domain specific vocabularies should be aligned with common vocabularies. The alignment process cannot always be done automatically as there are exceptions that require careful examination and manual work. Another focus is to improve scalability of the applications to anticipate problems that can occur in cases where the user is comparing very large collections. In addition to prototype improvements, we are planning to conduct a second in-depth evaluation with museology students. The goal of our evaluation is to investigate how comparison search applications can assist museology students learn about the various museum collections. We are interested in understanding the value of different features, such as the different visualizations, property selection and autocompletion for collection analysis task. Additionally, we want to assess the quality of our data alignments to support comparison search task.

6 Conclusion

LISA is a comparison search tool that enables comparison search across heterogeneous collections. It enables users to obtain an overview of characteristics of sets of cultural heritage objects and see the similarity and differences between them. Moreover, LISA enables users to examine collections from many different perspectives by using different visualizations and selecting among available distinguishing properties. We have implemented LISA with two cultural heritage dataset: RKD Image archive and Rijksmuseum Amsterdam, and conducted a preliminary evaluation with cultural heritage experts to collect early feedback on our system. In our evaluation, the cultural heritage experts indicated that comparison search tools can potentially be very helpful for students, researchers and curators. The experts also provided useful feedback to improve the system such as alignment quality and query formulation improvements.

Acknowledgments. We thank the participants from the Netherlands Institute for Cultural Heritage(ICN), the Digital Heritage Netherlands(DEN), and the Tropenmuseum Amsterdam, for providing us invaluable feedback on our application. We thank Hyowon Lee and Corrado Boscarino for their support in this research. This research was supported by the MultimediaN project funded through the BSIK programme of the Dutch Government and by the European Commission, contract FP6-027026, Knowledge Space of semantic inference for automatic annotation and retrieval of multimedia content K-Space.

References

1. S. Ahern, M. Naaman, R. Nair, and J. H.-I. Yang. World explorer: visualizing aggregate data from unstructured text in geo-referenced collections. In *JCDL '07*, pp 1–10, 2007.
2. A. Amin, L. Hardman, J. van Ossenbruggen, and A. van Nispen. Understanding cultural heritage experts' information seeking tasks. In *JCDL '08*, 2008.
3. A. Amin, M. Hildebrand, J. van Ossenbruggen, V. Evers, and L. Hardman. Organizing Suggestions In Autocompletion Interfaces. To appear in *ECIR '09*, 2009.
4. E. Callahan and J. Koenemann. A comparative usability evaluation of user interfaces for online product catalog. In *EC '00*, pp 197–206, 2000.
5. E. Hyvonen, et al. Culturesampo – finnish cultural heritage collections on the semantic web 2.0. In *DH-JAC-2009, Ritsumeikan, Kyoto, Japan*, 2009.
6. M. Hildebrand, J. R. van Ossenbruggen, A. K. Amin, L. Aroyo, J. Wielemaker, and L. Hardman. The Design Space Of A Configurable Autocompletion Component. cwi INS-E0708, CWI, 2007.
7. A. Kobsa. An empirical comparison of three commercial information visualization systems. In *INFOVIS '01*, pp 123, 2001.
8. J. Lee, H. S. Lee, and P. Wang. An interactive visual interface for online product catalogs. *Electronic Commerce Research*, 4(4):335–358, 2004.
9. T. Ni. Musa–a prototype for multiple-step aggregation visualization. In *CMV '05*, pp 15–21, 2005.
10. S. Perugini, K. McDevitt, R. Richardson, M. A. Pérez-Qui nones, R. Shen, N. Ramakrishnan, C. Williams, and E. A. Fox. Enhancing usability in citidel: multimodal, multilingual, and interactive visualization interfaces. In *JCDL '04*, pp 315–324, 2004.
11. C. Plaisant, J. Rose, B. Yu, L. Auvil, M. G. Kirschenbaum, M. N. Smith, T. Clement, and G. Lord. Exploring erotics in emily dickinson's correspondence with text mining and visual interfaces. In *JCDL '06*, pp 141–150, 2006.
12. G. Schreiber, et al. Semantic annotation and search of cultural-heritage collections: The multimedian e-culture demonstrator. *J. Web Sem.*, 6(4):243–249, 2008.
13. R. Shen, N. S. Vemuri, W. Fan, R. da S. Torres, and E. A. Fox. Exploring digital libraries: integrating browsing, searching, and visualization. In *JCDL '06*, pp 1–10, 2006.
14. M. Spenke, C. Beilken, and T. Berlage. Focus: the interactive table for product comparison and selection. In *UIST '96*, pp 41–50, 1996.
15. P. Steiger and M. Stolze. Effective product selection in electronic catalogs. In *CHI '97*, pp 291–292, 1997.
16. H. Stuckenschmidt and F. v. Harmelen. Knowledge-based validation, aggregation, and visualization of meta-data: Analyzing a web-based information system. In *WI '01*, pp 217–226, 2001.
17. T. Tenev and R. Rao. Managing multiple focal levels in table lens. In *INFOVIS '97)*, pp 59, 1997.
18. A. Tordai, B. Omelayenko, and G. Schreiber. Semantic excavation of the city of books. In *SAAKM*, 2007.
19. E. R. Tufte. *The visual display of quantitative information.* Graphics Press, Cheshire, CT, USA, 1986.
20. M. van Assem, V. Malaisé, A. Miles, and G. Schreiber. A method to convert thesauri to skos. pages 95–109, 2006.
21. J. Wielemaker, M. Hildebrand, J. Ossenbruggen, and G. Schreiber. Thesaurus-based search in large heterogeneous collections. In *ISWC '08*, pp 695–708, 2008.