

Chapter 1

Visualizing Search Results of Linked Open Data

C. Stab, D. Burkhardt, M. Breyer, K. Nazemi

Abstract Finding accurate information of high quality is still a challenging task particularly with regards to the increasing amount of resources in current information systems. This is especially true if policy decisions that impact humans, economy or environment are based on the demanded information. For improving search result generation and analyzing user queries more and more information retrieval systems utilize Linked Open Data and other semantic knowledge bases. Nevertheless, the semantic information that is used during search result generation mostly remains hidden from the users although it significantly supports users in understanding and assessing search results. The presented approach combines information visualizations with semantic information for offering visual feedback about the reasons the results were retrieved. It visually represents the semantic interpretation and the relation between query terms and search results to offer more transparency in search result generation and allows users to unambiguously assess the relevance of the retrieved resources for their individual search. The approach also supports the common search strategies by providing visual feedback for query refinement and enhancement. Besides the detailed description of the search system, an evaluation of the approach shows that the use of semantic information considerably supports users in assessment and decision-making tasks.

1.1 Introduction

Assessing information is a common task for decision makers. Especially in the area of policy modeling, analysts have to consider different perspectives when designing a new policy that impacts humans, economy and the environment. So homogeneous access possibilities as well as adequate representations for distributed data plays a

C. Stab, D. Burkhardt, M. Breyer, K. Nazemi
Fraunhofer Institute for Computer Graphics Research (IGD), Fraunhoferstr. 5, 64283 Darmstadt, Germany, e-mail: christian.stab@igd.fraunhofer.de

major role for providing adequate tools that facilitate decision making. Thereby, semantic technologies provide adequate tools for linking heterogeneous data sources and for generating broader contexts that facilitate information access and enables data exchange between different systems (Shadbolt et al., 2006). With the ongoing establishment of semantic technologies like the *Resource Description Framework* (RDF), the *Web Ontology Language* (OWL), semantic-oriented query languages like *SPARQL* and *Linked Open Data* (LOD) platforms like DBpedia (Auer et al., 2007) these developments are not only limited to specific domains but also adopted in daily search processes (Fernandez et al., 2008) and even for interlinking government data in so called *Linked Government Data* (LGD) platforms (Ding et al., 2010; Wood, 2011). For accessing these large amounts of interlinked data usually information retrieval methods are utilized that retrieve relevant resources by means of given search terms or keywords. The results of search processing are usually presented in sorted lists and in most cases the ordering of list entries represents the relevance of the results for the individual search according to various criteria (Cutrell et al., 2006). So the most relevant result is placed in the first row followed by less important ones. The semantic information of the resources that is used during the search result generation and the analysis of search terms remains hidden from the user in most cases, though this information considerably supports users in information-seeking tasks and selection of appropriate documents for further examination.

For designing search user interfaces efficient and informative feedback about the retrieved resources is critically important for the user to be able to assess the presented search results (Hearst, 2009). In particular this includes feedback about query formulation and about reasons the results were retrieved from the information system. However, the use of additional relevance indicators in result lists besides relevance ordering such as numerical scores or special icons turned out to be not successful for supporting users in understanding the retrieved results because the meaning of the relevance score is opaque to the user (White et al., 2007). This is because the majority of existing relevance indicators only presents a single relevance per search result that summarizes all criteria instead of offering a more fine-grained insight to search result processing.

In contrast to common search result presentations the presented approach makes use of information visualizations for representing search results of semantically modeled data. In order to offer users an adequate tool for assessing the relevance of the retrieved search results the approach combines information visualization techniques with semantic information and different weights that emerge during the search process. The approach also visually supports the common search strategies by providing visual feedback for query evolution. The main contributions and benefits of the presented approach are:

- *Query-Result-Relations*: The semantic search processing of the approach analyzes the given queries and identifies relations between query terms and search results. The visualization of these relations offers a fine-grained overview of search result relevancies and facilitates information seeking and assessment tasks.

- *Relevance Assessment*: The inclusion of semantic information in search result presentation offers more transparency in search result generation and successfully supports user in assessing the relevance of the presented resources.
- *Visual Query Enhancement*: The visual recommendation of additional query terms related to the set of search results supports the common search strategy and allows users to narrow current search results and to immediately receive visual feedback.

In the next section we give an overview of related work that influenced the development of our approach. Subsequently, we introduce our approach for processing and visualizing searches in semantic domains and give a detailed description of each component. We present the results of an evaluation comparing our approach to already existing solutions followed by an outlook of its application in the area of policy modeling.

1.2 Related Work

Semantic models are increasingly used for linking heterogeneous data sources as well as for generating broader contexts that facilitate information access (Shadbolt et al., 2006). Nowadays, these technologies are not only limited to specific domains but also adopted in daily search processes of web-based search engines (Fernandez et al., 2008). A commonly used and useful approach for representing search results is the term highlighting technique (Aula, 2004) where the terms of a given query are highlighted in the surrogates of search result lists. This approach is also referred to as Keyword-In-Context (KWIC). For example the BioText System (Hearst et al., 2007) represents beside extracted figures from relevant articles, query terms highlighted in the title and boldfaced in the text excerpts for communicating reasons the particular results were retrieved. Even though term highlighting can be useful for improving search result list presentations, it does not reveal the semantic interpretation of search results and prevent users from scanning the whole result list for getting an overview.

Although the initial intention of semantic technologies was not focused on presenting semantics to end-users, there are several approaches that benefit thereby. SemaPlorer (Schenk et al., 2009) is an interactive application that makes use of multiple semantic data sources and allows users to visualize results of their search in various views. The user interface of SemaPlorer combines a geographic visualization and a media view for visualizing geospatially annotated data and picture galleries respectively. The approach also includes facettation of search results and focuses on combining search results from different heterogeneous knowledge bases. Another approach of presenting semantic information to the user is the Relfinder interface (Heim et al., 2010). It supports users in interactively discovering relations between resources in semantic knowledge spaces. Users can prompt two or more

resources and the relations between them are shown in a graph-based visualization. Although this approach demonstrates the benefit of communicating semantic knowledge to users, it is strictly limited to relation discovery.

Other approaches utilize different information visualization techniques for improving search result presentations. To name only a few, the Microsoft Academic Search interface (Microsoft Academic Search, 2012) incorporates geographic, graph-based and temporal visualization techniques for searching and exploring publications or authors and offers also a stacked area chart for analyzing trends in the field of computer science. SkylineSearch (Stoyanovich et al., 2011) is a search interface that supports life science researchers in performing scientific literature search. It leverages semantic annotations to visualize search results in a scatterplot representing relevance against publication date. Even though semantic annotations are used for search processing and estimating relevance values, semantic knowledge is not directly presented to the user. The WebSearchViz (Nguyen and Zhang, 2006) is an approach for visualizing web search results based on the metaphor of the solar system. It offers users the possibility to observe the relevance between a query and a web search result by the spatial proximity between objects. However, the system does not visualize semantic interpretations of search results or semantic structures.

The visual design of the presented visualization is based on a force-based visualization method similar to RadVis (Hoffman et al., 1997). Force-based visualization methods utilize physical laws (e.g. Hooke's Law) to locate each data record on a two dimensional screen by assigning different forces between the visual representations of each record. By adapting these forces in an iterative simulation, the physical system reaches a mechanical equilibrium resulting in an aesthetical layout of the given data. Originally, RadVis is a visualization technique for multivariate data. For each dimension this visualization locates an anchor point on a circle. These anchor points can be seen as fixed ends of springs. Each representation of a data record is attached to all related anchor points and the attraction force for each spring is weighted with the value of the record for the specific dimension. This results in different distances between the data records and the anchor points that represent the characteristics of each record according to the visualized dimensions. This approach is especially useful for identifying outliers in the data and for recognizing clusters.

Although there are different approaches that make use of semantic information for improving search results or result presentation and approaches that utilize information visualization for representing search results, there is still a gap in combining semantic search processing, information visualization and search user interfaces. The approach presented in this chapter aims at combining these three technologies into an interactive search user interface that facilitates information access and relevance assessment.

1.3 Search Procedure

For visualizing search results, there are also additional requirements for the search procedure and in particular the result generation. In contrast to commonly used textual list presentations that are based on ordered result lists, for the visualization of search results structured data is needed. For instance these additional structures may include clusters, relations, or labeled taxonomies that can be exploited for providing meaningful visual representations. The semantic information provided e.g. in Linked Open Data databases provides a useful starting point for extracting this structural information.

The procedure used for retrieving the needed result structures from a semantic database consists of several steps at two different stages (Fig. 1.1). In the first stage the given user query is analyzed and dissembled into n-grams. Each of these n-grams is then mapped to an instance in the semantic data base e.g. by using an entity recognition approach like that described in (Paulheim and Fürnkranz, 2012). Additionally, the type label of each recognized entity is retrieved by identifying the most specific concept in the semantic structure. So the result of the query analysis is a set of n-grams related to an instance and its most specific concept. We also considered different filtering strategies for discarding incomplete results. In particular there are two different filters: (1) an entity filter that filters all n-grams that could not be mapped to an entity and (2) a type filter that removes all entities whose concepts could not be retrieved. The search system can be parameterized to use one, both or none of these filters.

In the second stage the results for each recognized entity are retrieved and extended with additional structural information. In the first step each recognized entity that meets the given filter condition is searched in the semantic data base using an indexed search. The result lists are combined using two sorting fields. The first sorting field refers to the number of occurrences for each retrieved resource in the retrieved lists and the second sort field is the number of references that each retrieved resource has in the underlying semantics data base. So the first element is the element that is retrieved for most of the recognized query terms. Based on the combined result list, the search system extracts relations between the recognized query terms and results and retrieves additional recommendations as resources related to the result list (Fig. 1.1). The relation extraction creates for each resource a set of relations to related n-grams. These relations are weighted according to the rank of the resource in the result list. Thus the result of the relation extraction is a set of relations between query parts that are recognized as entities and the results in the combined result list. For providing additional information about the search, the procedure retrieves also resources that are related to the result list. These recommendations are extracted by querying all related entities of the resources in the combined result list from the semantic database. For each of these related resources the most specific concept is extracted. The result is a set of concepts with associated resources that are related to the search result list. These concepts are ordered according to the number of related resources. So each recommendation is a concept with a set of associated resources that are related to the search result list.

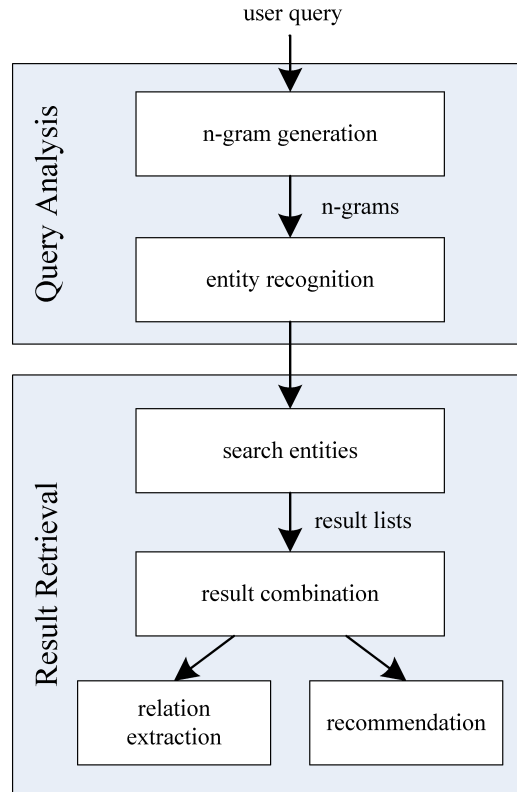


Fig. 1.1 Process for retrieving search results, query-result-relations and recommendations from semantic databases.

For our visualization approach we implemented the described search procedure as a web service using the DBpedia SPARQL-Endpoint. Besides the access to the semantics, the system also requires an indexed search for retrieving the result lists. For our prototype we used the DBpedia Lookup Service¹ that ranks the results according to the number of references (refCount). In the following section we describe how the results of the semantic search procedure are presented to the user and how the visualization interacts with the search system.

¹ for the prototype we use the DBpedia lookup service that is available at <http://wiki.dbpedia.org/lookup/>

1.4 Visualizing Semantic Search Results

Our visualization approach for representing search results in semantic domains is based on the semantics visualization framework SemaVis². The Framework includes several aspect-oriented visualization techniques (e.g. sunburst visualizations (Stab et al., 2010), timelines (Stab et al., 2010), map-based approaches (Nazemi et al., 2009), etc.) that can be combined to an application-specific visualization cockpit to represent different aspects of the underlying data. Several data providers for common semantic file formats and service APIs, a modular representation model for adapting the visual appearance as well as a script-based configuration language called *Semantics Visualization Markup Language* (SVML) are also included in the framework. In the following, we will first describe the basic concepts of the visualization component before introducing the details in succeeding sections.

The approach for visualizing semantic search results distinguishes between two different node types: (1) *Term Nodes* and (2) *Result Nodes*. Each node type is treated by different layout algorithms in the visualization and used to represent different information:

- *Term Nodes* represent the terms recognized by the described search process. These nodes are placed by a concentric layout algorithm at the startup of the visualization. Users are also able to freely move and order them on the surface according to their individual preferences.
- *Result Nodes* represent the hits in the combined result list that are found for the given user query. These nodes are visually connected to related term nodes with directed edges and are treated by a force-based layout algorithm according to their weights to the surrounding term nodes.

The placement of the result nodes in the center of the visualization is done by a force-based layout algorithm. The algorithm positions nodes in a two-dimensional space by assigning different forces to the edges and the nodes of a graph. These forces are adapted during the layout process in an iterative simulation until the physical system reaches a mechanical equilibrium. Due to this layout technique, the overlapping of nodes and edges is prevented as far as possible and an aesthetical layout of the graph is achieved. Another interesting characteristic of force-based layout methods is that the forces between the nodes can be weighted with different similarity values. As a result of the weighting, different distances between the nodes are derived and the placement of the nodes is affected. To exploit this feature we assign different values to the visualized result nodes and their edges and utilize a model based on the weights emerged during the semantic retrieval process. In the representation, a result node is then placed nearer to similar term nodes and further away from terms that are less related to it. So the placement of the result nodes in-

² SemaVis Framework: <http://www.semavis.net>

indicates the relevance of the results to the surrounding terms and users are able to distinguish different characteristics of retrieved hits.

Fig. 1.2 shows an example of the visualization approach³ for the query term '*apple computer steve jobs*'. The query analysis stage of the search process identified the six terms '*apple computer*', '*steve jobs*', '*apple*', '*computer*', '*steve*' and '*jobs*' and for 4 of these terms the most specific concepts '*company*', '*person*', '*plant*' and '*given name*'. Each of these identified terms is represented in a term node and the retrieved result nodes are attached to them using the weights identified during the relation extraction step. In the next section, we describe how this visualization metaphor can be used for providing more transparency in search processing and how it can be utilized for fostering search result comprehension.

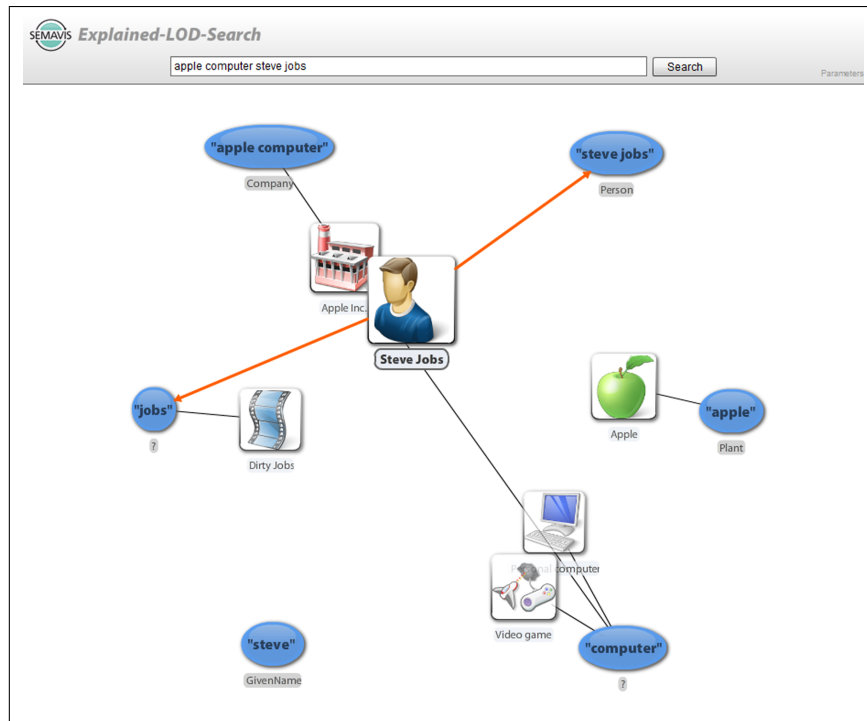


Fig. 1.2 Semantic search result visualization presenting the six most relevant search results for the query '*apple computer steve jobs*'. Each result is attached between the terms recognized during the semantic query analysis and according to the retrieved query-result-relations.

³ a version using the DBpedia database and the described search processing is available at <http://semanticsearch.semavis.net>

1.4.1 Query-Result-Relations

Semantic technologies enable information retrieval systems to 'understand' on the one hand the query terms given by the user (Sect. 1.3) and on the other hand different properties of the underlying resources. Based on the occurrences of the query terms either in the content, in the semantic properties or even in the semantic neighborhood, the underlying resources are filtered and ranked using information retrieval methods for presenting the result resource to the end user. In this process semantic information improves the retrieval process by providing an extended feature space for each resource as well as by offering methods for disambiguating and interpreting the query terms given by the user. However, the semantic information utilized during the retrieval process remains in most cases hidden from the user though it provides useful feedback about the reasons the results were retrieved. For example the query term '*ford*' might be interpreted as the name-property of a car manufacturer, as the surname-property of the famous inventor or the title-property of an activity for crossing rivers. Each of these interpretations will deliver a completely different result set. So it is not sufficient to only present the relations between recognized query terms and results, but it is also necessary to point out the semantic interpretation of the given query terms to allow an unambiguous assessment of retrieved results.

To meet these demands and to provide an adequate tool that allows users to unambiguously determine the most relevant result for their individual search, our approach visualizes both query-result-relations and the interpreted semantic meaning of query terms. Therefore, each term of the given query is presented in a term node of the visualization. The interpreted semantic meaning emerged during search processing (the label of the most specific concept) is visible in the label of the term node. So for every possible interpretation a new node is created that represents the query term and its retrieved interpretation. The relations between search results and the term nodes are depicted as directed and weighted edges between term nodes and result nodes. As mentioned above, the weighting of a query-result-relation is derived according to the rank of the result in the individual result list. As consequence of this weighting, the results are placed nearer to more relevant entities and term nodes respectively.

Fig. 1.3 shows two examples of the visualization approach representing two different queries and in each case the first ten hits of the result set. For the sake of clarity, we limited the number of simultaneously visible results to a fixed number and added common paging functions for switching between pages. The left example shows the visualization of the results for the query '*Henry Ford*' where the term '*Henry Ford*' is identified as person and the term '*Ford*' as company. Additionally, the term '*Henry*' is recognized by the query analysis step of the search processing but there is no specific type other than owl:Thing available. So the type for this query part is not evident. However, the visualization of the results reveals that only one of the results is related to the queried person whereas other results are related to the recognized company (Fig. 1.3 left).

The second example shows the visualization of the results for the query '*Pirates of the Caribbean*' (Fig. 1.3 right). For the given query the system identifies two

combination step of the described search process.

In order to make the optimum use of these values, each of them is mapped to specific visual properties like length, color and size that can be preattentively perceived (Ward et al., 2010). Thanks to the characteristic of preattentive perception, these relevance indicators can be perceived faster and easier by the users compared to common indicators that are often represented as textual percentages. To take these advantages, the result rank is used to adjust the size and the intensity of the result nodes. Thereby the resource that has the highest overall rank for a specific search query is presented most conspicuous whereas resources with minor rank are visualized less notable (Fig. 1.2). On the other hand the relation weights are used to adapt the weighting of edges between results and term nodes. This results in different lengths of the visible connections and indicates the relevance between specific query terms and search results.

1.4.3 Visual Support for Query Evolution

A search process of common users includes various search requests and queries until the needed information is found. Usually such a process starts with a general query that is revised in consecutive search queries until some resources for further exploration are found in the result set. This kind of search strategy was also revealed in several studies which showed that it is a common strategy for the user to first issue a general query, then review a few results, and if the desired information is not found, to reformulate or to enhance the query (Hearst, 2009) (Jansen et al., 2005) (Jansen et al., 2007). To support this behavior of common users when searching for information in query-based information retrieval systems, we integrated several features that enable users to interactively change their queries in the visualization.

In the presented visualization approach the query terms of the user are visualized in term nodes that are arranged in a circular form around the result nodes. According to this characteristic, the state of these visible term nodes and the included terms reflect the current search intention of the user. Transferred to the visualization approach, the before mentioned strategy of query evolution corresponds to a substitution, reassignment, creation or removal of term nodes. Hence, the change of the current state of the term nodes results in a new search condition that in turn changes the visible result set or the visible relations between result and term nodes. On the one hand, the creation of further terms defines a more specific search condition and on the other hand, the removal of term nodes results in wider-ranged search spaces. In contrast to commonly used search user interfaces, the influence of changing search conditions is immediately visible in the visualization. The representation of query-result-relations reveals which of the current search results fulfill new conditions (Fig. 1.4) and provides an immediate visual feedback of the users' query evolution.

- *Q1: With the help of the system, I was able to quickly and effectively solve the given task.*
- *Q2: The system presented the information needed for answering the question clear and unambiguous.*
- *Q3: Would you use the system in the future for similar search tasks?*

Each of these items was rated by the participants on a five point Likert scale from 1 (strongly disagree) to 5 (strongly agree).

1.5.1 Experimental Design

According to the hypothesis that contains one independent variable with two different conditions (list view and visualization) the design of our experiment is based on a basic design (Lazar et al., 2010). Additionally, we decided to use a within-group design for our experiment where each participant accomplishes the given tasks in each condition (in this case the different user interfaces). In contrast to between-group designed experiments, in within-group designs less participants are needed and individual differences between the participants are isolated more effectively (Lazar et al., 2010). Possible learn effects when switching between conditions are controlled by a systematic randomization of condition- and task-ordering. Furthermore, participants were advised to disregard the knowledge from previous conditions and to explicitly show the solution of tasks by means of elements in the user interface.

Altogether, the experiment contains three tasks that had to be accomplished from every participant with both conditions (list view and visualization). Because the focus of the evaluation is the comparison of two different user interfaces and not the investigation of the whole search process, we were able to pre-assign the query terms for every task. So every participant retrieves the same results for every task and thus also the same visual representation and the evaluation outcome is not influenced by other factors.

In the first task participants had to identify the relations between each search result and the recognized terms of the given query. The second task was of the same type as the first task with the difference that the result set contains more complex relations. In the third task participants had to identify the most relevant item for a specific search situation. To ensure that the solution could be found in each condition, we performed several pretests. We also ensured that each participant gets the same visual presentation for each task and condition. The time limit for each task was set to three minutes. If a wrong answer was given or a participant could not solve a task, the completion time of the task was also set to three minutes.

1.5.2 Procedure

17 participants, mainly graduates and students attended the evaluation. The participants were between 24 and 29 years old and mainly involved in computer science ($M = 4.65$; $SD = 0.6$)⁴. After a general introduction to the user study and an explanation of the procedure and tasks, participants got a brief introduction to both systems in systematically randomized ordering. Both systems were queried with a reference query and participants had the chance to ask questions about the systems. After each task, participants had to rate their overall satisfaction with the system on a scale from 1 to 9 and the three before mentioned items concerning their subjective opinion about the system on a Likert scale from 1 (strongly disagree) to 5 (strongly agree). After participants had completed all tasks, they had to answer a brief demographic questionnaire.

1.5.3 Results

Fig. 1.5 shows the average task completion times for each of the three tasks and both conditions. The direct comparison of the average task completion times reveals that participants performed better with our visualization approach ($avg(t) = 51.3$ sec; $SD = 25.8$) compared to the list view ($avg(t) = 88.1$ sec; $SD = 30.1$).




Fig. 1.5 Task completion time and user satisfaction

A paired-samples t-test also suggests that there is a significant difference in the task completion time between the group who used the list view and the group who used our visualization approach ($t(50)=7.8028$, $p < 0.05$). Hence, the null hypothesis is refuted and the alternative hypothesis confirmed. The comparison of means also indicates that users performed significantly faster with the visualization approach compared to the list presentation. So we can proceed from the assumption

⁴ Measured on a five point scale (5 = very much experience; 1 = very little experience) in the demographic part of the questionnaire.

that visualizing search results taking semantic information into account has a positive effect on the efficiency when assessing the relevance of search results.

Table 1.1 Results of the subjective ratings indicate that users prefer the visualization approach.

Presentation	Q1	Q2	Q3
Visualization	4.33	4.37	4.25
Listview	3.15	3.03	3.14

The evaluation of satisfaction ratings indicates that participants feel more comfortable with our visualization approach instead of the commonly used list view. The list view obtained an average rating of 5.31 with a standard deviation of 1.91 whereas the visualization obtained an average rating of 7.57 and a standard deviation of 1.10. Additionally, the result of the subjective ratings (Table 1.1) and in particular question “*Would you use the system in the future for similar search tasks?*” confirms the assumption that users prefer the visualization to the list presentation (list: $M = 3.14$; $SD = 0.87$; visualization: $M = 4.25$; $SD = 0.77$).

1.6 Visualizations and Linked Data in the Policy Modeling Process

Retrieving and accessing information is a challenging task and crucially important in many different domains. This is especially true in the area of policy modeling where decisions impact humans, economy or environment. The creation of novel policies is a very complex task that requires several process steps (Macintosh, 2004) to ensure validity and positive effects. It is easy to imagine that an insufficient analysis of the underlying problem and the consideration of all impact factors will result in a policy that fails the intended goals. Accurate decision making in this domain not only requires the consideration of diverse impact factors but also the inclusion of increasingly complex and dynamic scenarios. For improving this process and the quality of the achieved policies respectively, recent initiatives, like Open Government Data or Linked Government Data aim at publishing and interlinking vast amounts of data for enabling accurate decision support and innovative ICT solutions for fostering political decision making. However, the amount of available data that contains implicit and hidden information relevant for specific policy decisions and scenarios continuously increases and the access gets more and more complicated. To improve the retrieval tools, not only the interlinking of open administrative data gains enormous importance for policy modeling but also the development of novel result presentations and exploration tools. Present systems for searching and accessing this information are currently limited to textual result presentations and require comprehensive knowledge about the domain for finding the information that fits to the specific case. Approaches, like the presented visualization technique, will on

the one hand provide a homogenous access for distributed and interlinked data and on the other hand enable political decision makers to identify unknown and hidden information. In particular during the information foraging step of the policy modeling process (Kohlhammer et al., 2012), visualization techniques will enable an optimal analysis of the need for a policy and accurate assessment of issues relevant to a specific scenario. In order to further develop this idea, we investigate different visualization approaches along each step of the policy modeling process in the European project FUPOL that proposes a comprehensive and new governance model to support the policy design lifecycle.

1.7 Conclusion & Future Work

In this chapter, we introduced an approach for visualizing search results of Linked Open Data. The results of the evaluation showed that the utilization of semantic information in search results visualization successfully fosters search result comprehension and supports users in assessing retrieved resources. Also the approach performed well for presenting different semantic interpretations of query terms and query-result-relations respectively. The visual recommendation of related items and immediate visual feedback for query refinement additionally fosters the common search strategies of users and offers more transparency in search result processing. The chapter also outlined the application of the approach in the domain of Policy Modeling and presented some of the advantages of information visualizations in this domain.

For the future work it might be interesting to include features of statistics data in the retrieval process and in particular in the attribute set of the proposed visualization approach. Beside the possibility of searching in the content of statistics this feature might also improve the search and the visual query refinement feature.

References

- Auer S, Bizer C, Lehmann J, Kobilarov G, Cyganiak R, Ives Z (2007) DBpedia: A Nucleus for a Web of Open Data. In: Proceedings of the 6th International and 2nd Asian Semantic Web Conference (ISWC2007+ASWC2007), pp. 715-728
- Aula A (2004) Enhancing the readability of Search Result Summaries. In: Proceedings of HCI 2004, pp. 6-10
- Berners-Lee T (2006) Linked Data. <http://www.w3.org/DesignIssues/LinkedData.html> Accessed 17 July 2012
- Bizer C, Heath T, Berners-Lee T (2009) Linked Data - The Story So Far. In Int. Journal on Semantic Web and Information Systems (JSWIS)
- Cutrell E, Robbins D, Dumais S, Sarin R (2006) Fast, Flexible Filtering with Phlat. In: Proceedings of the SIGCHI conference on Human Factors in computing systems, pp. 261-270

- Ding L, DiFranzo D, Graves A, Michaelis J, Li X, McGuinness DL, Hendler JA (2010) TWC Data-Gov Corpus: Incrementally Generating Linked Government Data from data.gov. In: International World Wide Web conference, WWW 2010
- Fernandez M, Lopez V, Sabou M, Uren V, Vallet D, Motta E, Castells P (2008) Semantic Search Meets the Web. In: 2008 IEEE International Conference on Semantic Computing, pp. 253-260
- Hearst MA (2009) Search User Interfaces. Cambridge University Press
- Hearst MA, Divoli H, Guturu A, Ksikes P, Nakov MA, Wooldridge JY (2007) BioText Search Engine: Beyond Abstract Search. In: Bioinformatics, 23(16):2196
- Heim P, Lohmann S, Stegemann T (2010) Interactive Relationship Discovery via the Semantic Web. In: Proceedings of the 7th Extended Semantic Web Conference (ESWC 2010), LNCS 6008, pp. 303-317
- Hoffman P, Grinstein G, Marx K, Grosse I, Stanley E (1997) DNA Visual and Analytic Data Mining. In: Proceedings of Visualization '97, pp. 437-441
- Jansen, BJ, Spink, A, Koshman, S (2007) Web Searcher Interaction with the Dogpile.com Metasearch Engine. Journal of the American Society for Information Science and Technology, 58(5), pp. 744-755
- Jansen BJ, Spink A, Pedersen, JO (2005) A Temporal Comparison of Altavista Web Searching. Journal of the American Society for Information Science and Technology, 56(6), pp. 559-570
- Kohlhammer J, Nazemi K, Ruppert T, Burkhardt D (2012) Towards Visualization in Policy Modeling. Computer Graphics and Applications, IEEE, Vol. 32(5), pp. 84-89
- Lazar J, Feng JH, Hochheiser H (2010) Research Methods in Human - Computer Interaction. John Wiley & Sons
- Macintosh A (2004) Characterizing E-Participation in Policy-Making. In: Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04), Vol. 5. IEEE Computer Society, Washington, DC, USA
- Microsoft Academic Search (2012) <http://academic.research.microsoft.com>. Accessed 17 July 2012
- Nazemi K, Breyer M, Hornung C (2009) SeMap: A Concept for the Visualization of Semantics as Maps. In: Proceedings of the 5th International Conference on Universal Access in Human-Computer Interaction (UAHCI '09). Part III, Springer, pp. 83-91
- Nguyen T, Zhang JA (2006) Novel Visualization Model for Web Search Results. IEEE Transactions on Visualization and Computer Graphics. 12(5), pp. 981-988
- Paulheim H, Fürnkranz J: Unsupervised Generation of Data Mining Features from Linked Open Data, In: International Conference on Web Intelligence and Semantics (WIMS'12), pp. 31:1-31:12
- Schenk S, Saathoff C, Staab S, Scherp A (2009) SemaPlorer - Interactive Semantic Exploration. In: Journal of Web Semantics: Science, Services and Agents on the World Wide Web, 7(4), pp. 298-304
- Shadbolt N, Berners-Lee T, Hall W (2006) The Semantic Web Revisited. In: IEEE Intelligent Systems 21(3), pp. 96-101
- Stab C, Breyer M, Nazemi K, Burkhardt D, Hofmann C, Fellner DW (2010) SemaSun: Visualization of Semantic Knowledge based on an improved Sunburst Visualization Metaphor. In: Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2010 pp. 911-919. Chesapeake, VA: AACE
- Stab C, Nazemi N, Fellner DW (2010) SemaTime - Timeline Visualization of Time-Dependent Relations and Semantics. Advances in Visual Computing - 6th International Symposium, ISVC 2010, pp. 514-523
- Stoyanovich J, Lodha M, Mee W, Ross KA (2011) SkylineSearch: Semantic Ranking and Result Visualization for Pubmed. In: Proceedings of the 2011 International Conference on Management of Data, SIGMOD '11, pp. 1247-1250
- Ward M, Grinstein G, Keim D (2010) Interactive Data Visualization: Foundations, Techniques, and Applications. A. K. Peters, Ltd., Natick

- White RW, Bilenco M, Cucerzan S (2007) Studying the Use of Popular Destinations to Enhance Web Search Interaction. In: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 159-166
- Wood D (2011) Linking Government Data. Springer New York Dordrecht Heidelberg London