

Semantics Visualization for Fostering Search Result Comprehension

Christian Stab, Kawa Nazemi, Matthias Breyer, Dirk Burkhardt,
and Jörn Kohlhammer

Fraunhofer Institute for Computer Graphics Research IGD
64283 Darmstadt, Germany

{christian.stab,kawa.nazemi,matthias.breyer,dirk.burkhardt,
joern.kohlhammer}@igd.fraunhofer.de

Abstract. Current search engines present search results in an ordered list even if semantic technologies are used for analyzing user queries and the document contents. The semantic information that is used during the search result generation mostly remains hidden from the user although it significantly supports users in understanding why search results are considered as relevant for their individual query. The approach presented in this paper utilizes visualization techniques for offering visual feedback about the reasons the results were retrieved. It represents the semantic neighborhood of search results, the relations between results and query terms as well as the relevance of search results and the semantic interpretation of query terms for fostering search result comprehension. It also provides visual feedback for query enhancement. Therefore, not only the search results are visualized but also further information that occurs during the search processing is used to improve the visual presentation and to offer more transparency in search result generation. The results of an evaluation in a real application scenario show that the presented approach considerably supports users in assessment and decision-making tasks and alleviates information seeking in digital semantic knowledge bases.

Keywords: Semantic Search, Information Visualization, SemaVis, Search User Interface, Visual Query Enhancement.

1 Introduction

The optimal use of information and knowledge plays a major role in global competition and forms the basis for competitiveness of industrial companies. Thereby, semantic technologies provide adequate linking tools for heterogeneous data sources as well as the generation of a broader context that facilitates information access and enables data exchange between different systems [1]. With the ongoing establishment of semantic technologies like the *Resource Description Framework (RDF)*¹, the *Web Ontology Language (OWL)*² and semantic-oriented query languages

¹ <http://www.w3.org/RDF/>

² <http://www.w3.org/TR/owl-features/>

like *SPARQL*³ these developments are not only limited to specific domains but also adopted in daily search processes of web-based search engines [2]. In both, domain-specific applications and web-based search engines, the results of search processing are usually presented in sorted lists. In most cases the ordering of list entries represents the relevance of the results for the individual search of the user according to various criteria [3]. So the most relevant result is placed in the first row followed by less important ones. Using this kind of result presentation, the semantic information of the documents that is used during search result generation and the analysis of search terms, remains in most cases hidden from the user, though this information considerably supports users in information-seeking tasks and selection of appropriate documents for further examination.

According to Hearst [4] efficient and informative feedback is critically important for designing search user interfaces. This includes in particular feedback about query formulation and about reasons the particular results were retrieved. However, relevance indicators besides list ordering such as numerical scores or special icons are less frequently used because the meaning of the relevance score is opaque to the user [5] in these presentations. 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 order to offer users an adequate tool that provides nevertheless the possibility to assess the relevance of retrieved search results, we developed a novel approach that utilizes information visualization techniques and semantic information that emerges during search result generation. The major contributions and benefits of our approach are:

- *Support for relevance assessment:* The presented approach supports users in assessing the relevance of search results and offers more transparency in the search result generation process.
- *Query-Result-Relation visualization:* The visual representation of relations between query terms and search results as well as the retrieved semantic meaning of query terms offers a fine-grained visual overview of search result relevancies and facilitates the information seeking and decision making process.
- *Visual feedback for query-enhancement:* The illustration of additional attributes and possible terms related to a given search request allows users to narrow search results and to refine the individual search process.

The rest of the paper is organized as follows: In the next section we introduce our approach for presenting search results in semantic domains and give a detailed description of all parts and features. Then we introduce the application scenario of the visualization and give an overview of its domain. We present the evaluation that we have performed to compare our approach to already existing solutions followed by a related work section, a discussion and a prospect of future work. As a detailed description of the whole search process with all technical aspects is beyond the scope of this paper, we only briefly describe the semantic background processing and focus on the aspects of the visualization component and the advantages of semantics for visualizing search results.

³ <http://www.w3.org/TR/rdf-sparql-protocol/>

2 Visualizing Search Results in Semantic Domains

The visualization approach is part of the SemaVis-Framework⁴, an adaptive visualization framework that contains different aspect-oriented visualizations [6, 7, 8] for presenting semantic information. The framework also provides different data providers, modular presentation models and a script-based configuration language that facilitate the development and integration of novel visualization techniques in semantically modeled environments. In the following we will give an overview of the visualization component before introducing the details in succeeding sections.

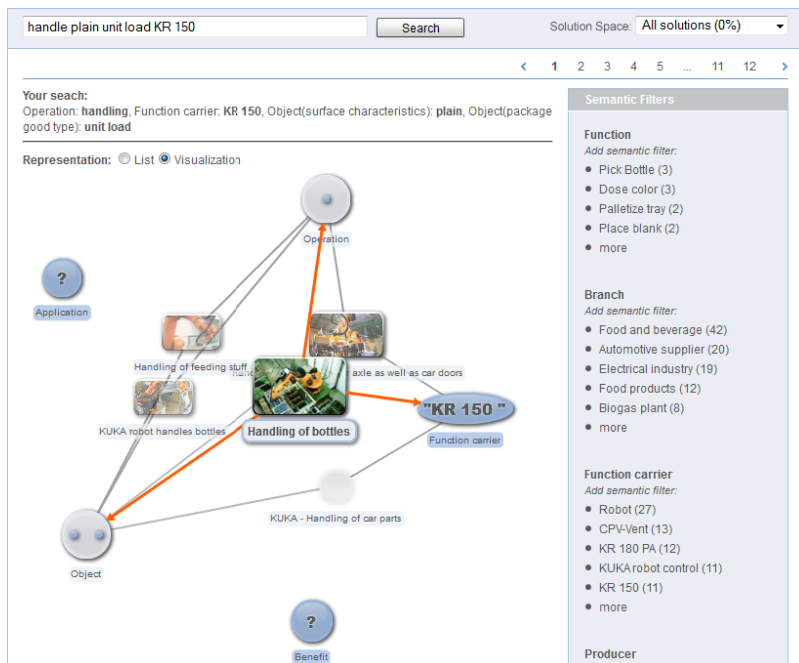


Fig. 1. Semantics visualization of search results showing query-result-relations, retrieved semantic meanings of query terms and result relevancy for an exemplary search

The visualization component is mainly based on a force-based layout algorithm that positions nodes in a two-dimensional space by assigning different forces to the edges and nodes of the graph. By adapting these forces in an iterative simulation, the physical system reaches a mechanical equilibrium resulting in an aesthetical layout of the graph. In contrast to commonly used force-based algorithms that assign the same force to every node and edge respectively, our approach utilizes a weighted model

⁴ <http://www.semavis.com>

based on similarities emerged during the semantic retrieval process. Another difference is that the visualization distinguishes between different node types:

- *Attribute nodes* are not positioned by the force-based algorithm. Instead, they are placed by a second layout algorithm in a circular form during the initialization of the visualization and users are allowed to freely move them on the surface (Fig. 1 shows an example).
- *Result nodes* represent the hits found for the given search query. They are suspended between the attribute nodes and are positioned by the force-based layout algorithm according to their similarities and relations to the attribute nodes.

Figure 1 shows an example with five attribute nodes and five result nodes. The weighting of the attraction and repulsion of the nodes and edges according to the retrieved similarity values arranges the result nodes closer to more similar attributes. So the best hit that matches to all attributes is placed near the center of the visualization.

2.1 Visualizing Query-Result-Relations

Giving adequate feedback about the reasons the results were retrieved is one of the major challenges for designing adequate search user interfaces. This is especially important for semantic search engines, in which the meaning of query terms is interpreted by means of semantically modeled entities, because the interpretation might be highly ambiguous. For example the query term *ford* might be interpreted as the name-attribute of a car manufacturer, as the surname-attribute of the famous inventor or the title-attribute 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 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 an attribute node of the visualization. The interpreted semantic meaning emerged during search processing is visible in the label of the attribute node. So for every possible interpretation a new node is created that represents the query term and its semantic meaning. The relations between search results and the instantiated⁵ attribute nodes are depicted as directed and weighted edges between attribute nodes and result nodes. As mentioned above, the weighting of an edge is derived from the retrieved similarity between the result and the attribute node, whereby the results are placed nearer to more relevant query terms and attribute nodes respectively.

⁵ 'instantiated' in this context means that a query term is assigned to a specific attribute.

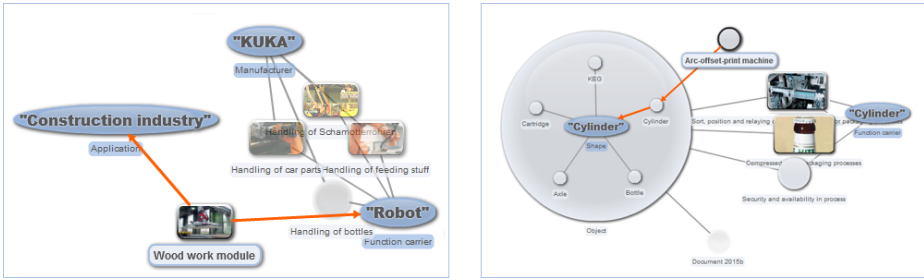


Fig. 2. Left: The visualization of query-result-relations reveals that only one of the five results is semantically related to the queried application area. Right: The visual representation of the identified semantic meanings of query terms avoids mistakes and ambiguity in result assessment tasks.

Figure 2 shows the result visualization of the query ‘*kuka robots in construction industry*’, where the term *kuka* is identified as manufacturer, the term *robot* as function carrier and *construction industry* as application area. The visualization reveals that only one of the results is related to the queried application area whereas other results are related to the given manufacturer (Figure 2 left). The second example shows the visualization of the results for the query ‘*cylinder*’. The given term is on the one hand identified as shape of an object and on the other hand as a specific function carrier. By visualizing the connections between search results and related interpretations of the query term, users can easily recognize the results that match their initial search intention.

2.2 Semantic Neighborhood and Hierarchical Attributes

Search results in semantic domains are not only retrieved by analyzing the content of resources but also by considering the semantic information and the semantic structure respectively. For example a resource that matches to only one of the given query terms is higher rated in the result list when the remaining terms match to semantically related resource. In some cases, semantic search processing enables the retrieval of highly relevant resources even if the given query terms are not contained in the resources. Especially in such cases where semantic structures are responsible for result generation, it can be a very time-consuming and tedious task to identify the right results for the individual search process. So it is important to provide an adequate presentation that allows users to unambiguously assess the retrieved results and enables them to comprehend why specific results are considered as relevant. To offer this kind of feedback the proposed approach presents related resources that are responsible for result retrieval and their related attributes in *expandable attribute nodes*. Thus each of these nodes contains resources from the semantic neighborhood of retrieved results that are of some relevance for the result generation. The labels of these expandable attribute nodes are derived from the conjoint concept in the semantic structure to indicate their meaning.

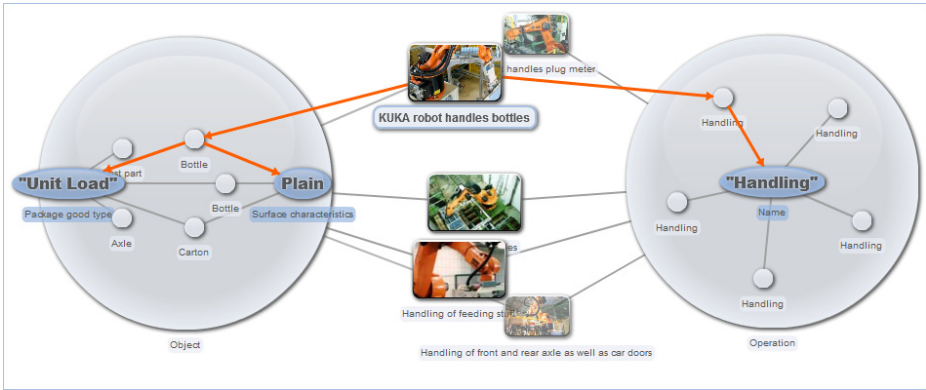


Fig. 3. The visualization of the semantic neighborhood combined with query-result-relations reveals why search results are retrieved during semantic search processing

Figure 3 shows an example of two hierarchically structured attribute nodes. The expandable attribute node on the left side contains five different objects (bottle, carton, etc.) and two elementary attribute nodes that are instantiated with given query terms. Each of the five object nodes is suspended between the inner attribute nodes to display the relations and relevancies to the current search. Finally, the relations between search result and query terms are represented by highlighting the path from a result over the related resource up to the instantiated attribute nodes. So the visualization indicates the ‘indirection’ in search result processing and reveals parts of the semantic neighborhood that are responsible for search result retrieval.

2.3 Mapping Results’ Relevance to Visual Properties

Beside the visualization of relations between query terms and retrieved results the proposed approach utilizes different similarity values that emerge during the result retrieval process for improving the visualization:

- *Partial similarities* are a measure between attribute nodes and results that represent the relevance of a retrieved resource to a given query term (encapsulated in an attribute node).
- *Result similarities* are aggregated values of all partial similarities that correspond to the overall relevance of a retrieved result.

In order to make the optimum use of these values, each similarity is mapped to specific visual properties like length, color and size that can be preattentively perceived [9]. On the one hand the size and color intensity of result nodes are adjusted according to the result similarity. Thereby the resource that has the highest overall similarity for a specific search query is presented most conspicuous whereas resources with minor similarities are visualized less notable (Figure 1). On the other hand partial similarities are used to adapt the weights of edges between results and attribute

nodes. This results in different lengths of the visible connections and indicates the relevance between specific query terms and search results.

2.4 Visual Feedback for Query Enhancement

Several studies revealed that it is a common search 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 [4, 10, 11]. Transferred to the presented visualization, this refinement strategy corresponds to substitutions or reassignments of attribute values because these are directly related to the terms of the current search query. On the one hand the instantiation of further attributes defines a more specific search condition and on the other hand the removal of attribute values 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 (Figure 4) and provides an immediate visual feedback for the query refinement task.

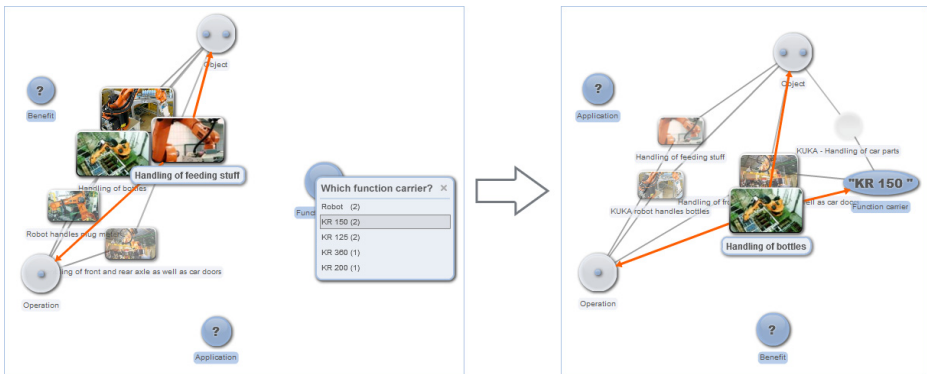


Fig. 4. The visual recommendation of additional attributes and possible terms for query enhancement offers a visual tool for narrowing search results

To ensure that users are aware of additional attributes the visualization recommends attributes that are not instantiated by the given query but related with the current result set. These recommendations are visualized as additional attribute nodes and labeled with a question mark to encourage users to instantiate them for narrowing their search. The size of the recommended attribute nodes is mapped to their influence to the current result set. So attribute nodes whose instantiation will cause major changes of the result set are represented larger than attribute nodes whose instantiation will only affect smaller parts. By selecting a specific recommendation, users are able to select different values for instantiating the attribute node and narrowing their retrieved results (Figure 4).

3 Application Scenario

The developed visualization approach is applied in the field of mechanical engineering and automation technology where highly complex processes and diverse user groups are involved. The processes in these domains range from initial development over construction and production steps up to sales and consumer services. So the same resource is treated in different contexts and various software systems. To ensure that each employee in these complex processes has access to the appropriate information at the right time, semantic technologies are used to link different knowledge bases and to provide a novel way to access information.

In this context the objective of the visualization is to provide a homogeneous access to the combined knowledge base for a variety of users. The focus of the investigation is predominantly on providing more transparency in search processing and to offer a tool that enables users to unambiguously assess the results of individual search processes. Additionally the aspect of recommending further dimensions for improving and narrowing the result set plays a major role in the investigation of the introduced visualization tool. In the current state the visualization is fully integrated in the search platform⁶.

4 Evaluation

For evaluating our approach we performed a user study in which we compared the visualization with a common list presentation (Figure 5). The study is mainly focused on answering the question whether our visualization approach can support the user in assessing search results and if our approach satisfies the needs of searchers. For verification of our assumption we investigated the task completion time and formulated the following hypothesis:

- *H1: There is a difference in task completion time between the list presentation and the visualization in assessing search results.*

Additionally to the task completion time we measured the user satisfaction as a subjective evaluation criterion.

4.1 Experimental Design

According to the hypothesis that contains one independent variable with two different conditions (list presentation and visualization) the design of our experiment is based on a basic design [12]. Additionally, we decided to use a within-group design for our experiment where each participant accomplishes the given tasks in each

⁶ Demonstration is available at <http://athena.igd.fraunhofer.de/Processus/semavis.html> Note that the knowledgebase of the online demonstrator is currently only available in German and contains only selected resources. Possible queries for demonstration are 'kuka roboter bauindustrie', 'glattes stückgut handhaben' and 'glas transportieren'.

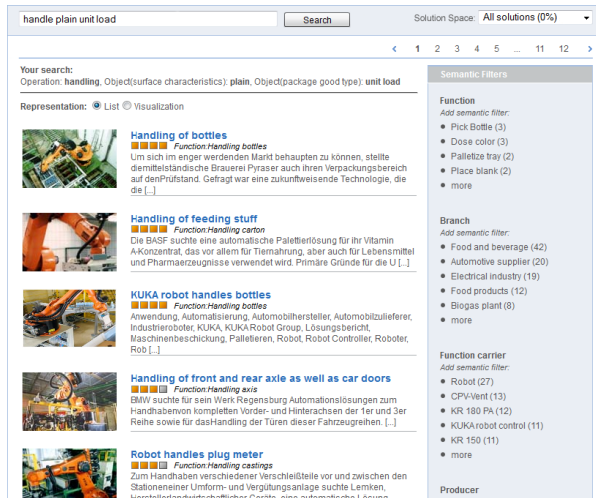


Fig. 5. List presentation of search results used for the evaluation

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 [12]. 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 presentation 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 terms of the given query. The second task was of the same type as the first task with the difference that the result contains hierarchical structured attributes instead of only flat attributes. 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.

4.2 Procedure

Altogether 17 participants, mainly graduates and students attended the evaluation. The average participant was between 24 and 29 years old. The participants were mainly involved in computer science ($M = 4.65$; $SD = 0.6$)⁷ and had no previous knowledge of the engineering domain. 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 three additional questions concerning their subjective opinion of 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.

4.3 Results

Figure 6 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 presentation ($avg(t) = 88.1$ sec; $SD = 30.1$). 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 presentation and the group who used our visualization approach ($t(50)=7.8028$, $p<0.05$).

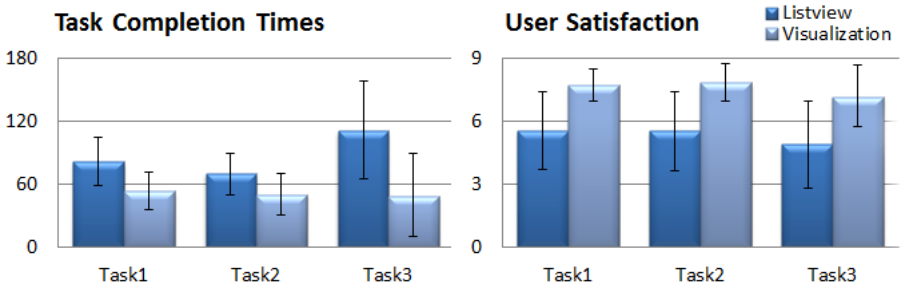


Fig. 6. Left: Task completion times. Right: Average user satisfaction

Hence the null hypothesis is refuted and the alternative hypotheses 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 that visualizing search results taking semantic information into account has a positive effect on the efficiency when assessing search result relevance.

⁷ Measured on a five point scale (5 = very much experience; 1 = very little experience) in the demographic part of the questionnaire.

The evaluation of satisfaction ratings indicates that participants feel more comfortable with our visualization approach instead of the commonly used list presentation. The list visualization 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 question “Would you use the system in the future for similar searching tasks?” confirms the assumption that users prefer the visualization to the list presentation (list: $M = 3.14$; $SD = 0.87$; visualization: $M = 4.25$; $SD = 0.77$)⁸.

5 Related Work

Although the objective of semantic technologies was not focused on presenting semantics to end-users, there are several other approaches that benefit thereby. SemaPlorer [13] is an interactive application that allows users to visualize the search results of multiple semantics data sources. The user interface of SemaPlorer also provides a geographic visualization and a media view for visualizing geospatially annotated data and picture galleries respectively. However, this approach is mainly focused on combining search results from different heterogeneous knowledge bases and faceting the search by predefined facets. The Relfinder interface [14] 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 between two or more resources.

There are also different approaches for using information visualization techniques for search user interfaces. To name only a few, the Microsoft Academic Search interface [15] incorporates geographic, graph-based and temporal visualization techniques for exploring publications or authors and offers also an stacked area chart for analyzing trends in the field of computer science. SkylineSearch [16] 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 plotting 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 [17] is an approach for visualizing web search results based on the metaphor of the solar system. It offers users the possibility to observe the semantic relevance between a query and a web search result by the spatial proximity and distance between objects. However the system does not visualize semantic interpretations of search results or semantic structures.

Another commonly used and useful approach for visualizing result relevancy is the term highlighting technique [18] where the terms of the given query are highlighted in the surrogates of search result lists. For example the BioText System [19] represents

⁸ Measured on a five point Likert scale.

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.

6 Discussion

The introduced approach was applied and evaluated in the field of mechanical engineering and automation technology. Although this domain contains highly complex processes and different kinds of heterogeneous users, domain experts were able to semantically design it and build a comprehensive model that enables different stakeholders the access to heterogeneous resources. In such well-defined domains, aspects like data diversity, user roles and processes are in some way controllable and the data access methods can be accurately aligned to specific tasks of the stakeholder. The results of the evaluation showed that the proposed visualization approach performed very well in the present domain. Nevertheless, further investigations are needed to prove if the proposed approach is also transferable to other domains and if it can be seamlessly integrated in semantic web search engines.

Currently, most search user interfaces are based on result list presentations and usually show the titles and surrogates of the results. Cause of the public's great familiarity with this commonly used search result presentation, there is a certain degree of risk with the introduction of a novel approach in user interfaces. Even if novel approaches provide a variety of extended features and easier information access, the success of each innovation in user interfaces is measured by the acceptance of the users. Although the results of the evaluation show that the introduced visualization approach performed well in a controlled experimental environment and users are convinced of its benefits, there is still the need to prove if visualization techniques will be applicable in web search engines. However, current trends show an increased use of information visualization techniques in search user interfaces.

7 Conclusion and Future Work

In this paper we introduced a novel approach for visualizing search results in semantic knowledge bases. The results of the evaluation showed that the utilization of semantic information in search results visualization successfully fosters search result comprehension and supports user 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 novel dimensions and immediate visual feedback for query refinement additionally fosters the common search strategies of users and offers more transparency in search result processing.

For future work we plan the extension of query refinement features. In particular we plan to implement the removal and change of attribute values that is not included

in the current version. Furthermore the multiple instantiation of attributes may be a useful extension of the introduced concept.

Acknowledgements. This work has been carried out within the Core-Technology Cluster (Innovative User Interfaces and Visualizations) of the THESEUS research program, partially funded by the German Federal Ministry of Economics and Technology. We thank H. J. Hesse, R. Traphöner and C. Dein (THESEUS PROCESSUS, Attensity Europe GmbH) for the inspiring discussions, the provision of the data and the support during the development of the data connection. We are also grateful to all participants that spend their time participating in the evaluation.

References

1. Shadbolt, N., Berners-Lee, T., Hall, W.: The Semantic Web Revisited. *IEEE Intelligent Systems* 21(3), 96–101 (2006)
2. Fernandez, M., Lopez, V., Sabou, M., Uren, V., Vallet, D., Motta, E., Castells, P.: Semantic Search Meets the Web. In: 2008 IEEE International Conference on Semantic Computing, pp. 253–260 (2008)
3. Cutrell, E., Robbins, D., Dumais, S., Sarin, R.: Fast, flexible filtering with phlat. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 261–270 (2006)
4. Hearst, M.A.: *Search User Interfaces*. Cambridge University Press (2009)
5. White, R.W., Bilenko, M., Cucerzan, S.: 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 (2007)
6. Stab, C., Breyer, M., Nazemi, K., Burkhardt, D., Hofmann, C., Fellner, D.W.: 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. AACE, Chesapeake (2010)
7. Stab, C., Nazemi, K., Fellner, D.W.: SemaTime - Timeline Visualization of Time-Dependent Relations and Semantics. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Chung, R., Hammound, R., Hussain, M., Kar-Han, T., Crawfis, R., Thalmann, D., Kao, D., Avila, L. (eds.) ISVC 2010, Part III. LNCS, vol. 6455, pp. 514–523. Springer, Heidelberg (2010)
8. Nazemi, K., Breyer, M., Hornung, C.: SeMap: A Concept for the Visualization of Semantics as Maps. In: Stephanidis, C. (ed.) UAHCI 2009, Part III. LNCS, vol. 5616, pp. 83–91. Springer, Heidelberg (2009)
9. Ward, M., Grinstein, G., Keim, D.: *Interactive Data Visualization: Foundations, Techniques, and Applications*. A. K. Peters, Ltd., Natick (2010)
10. Jansen, B.J., Spink, A., Pedersen, J.O.: A Temporal Comparison of Altavista Web Searching. *Journal of the American Society for Information Science and Technology* 56(6), 559–570 (2005)
11. Jansen, B.J., Spink, A., Koshman, S.: Web Searcher Interaction with the Dogpile.com Metasearch Engine. *Journal of the American Society for Information Science and Technology* 58(5), 744–755 (2007)

12. Lazar, J., Feng, J.H., Hochheiser, H.: *Research Methods in Human - Computer Interaction*. John Wiley & Sons (2010)
13. Schenk, S., Saathoff, C., Staab, S., Scherp, A.: SemaPlorer - Interactive Semantic Exploration. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web* 7(4), 298–304 (2009)
14. Heim, P., Lohmann, S., Stegemann, T.: Interactive Relationship Discovery via the Semantic Web. In: Aroyo, L., Antoniou, G., Hyvönen, E., ten Teije, A., Stuckenschmidt, H., Cabral, L., Tudorache, T. (eds.) *ESWC 2010, Part I. LNCS*, vol. 6088, pp. 303–317. Springer, Heidelberg (2010)
15. Microsoft Academic Search, <http://academic.research.microsoft.com>
16. Stoyanovich, J., Lodha, M., Mee, W., Ross, K.A.: SkylineSearch: semantic ranking and result visualization for pubmed. In: *Proceedings of the 2011 International Conference on Management of Data, SIGMOD 2011*, pp. 1247–1250 (2011)
17. Nguyen, T., Zhang, J.: A Novel Visualization Model for Web Search Results. *IEEE Transactions on Visualization and Computer Graphics* 12(5), 981–988 (2006)
18. Aula, A.: Enhancing the readability of search result summaries. In: *Proceedings of HCI 2004*, pp. 6–10 (2004)
19. Hearst, M.A., Divoli, H., Guturu, A., Ksikes, P., Nakov, M.A., Wooldridge, J.Y.: BioText Search Engine: beyond abstract search. *Bioinformatics* 23(16), 2196 (2007)