

Knowminer Search - a Multi-Visualisation Collaborative Approach to Search Result Analysis

Manuela Rauch*, Werner Klieber*, Ralph Wozelka*, Santokh Singh* and Vedran Sabol*[†]

*Know-Center GmbH

Graz, Austria

Email: {mrauch,wklieber,rwozelka,ssingh,vsabol}@know-center.at

[†]Graz University of Technology

Graz, Austria

Abstract—The amount of information available on the internet and within enterprises has reached an incredible dimension. Efficiently finding and understanding information and thereby saving resources remains one of the major challenges in our daily work. Powerful text analysis methods, a scalable faceted retrieval engine and a well-designed interactive user interface are required to address the problem. Besides providing means for drilling-down to the relevant piece of information, a part of the challenge arises from the need of analysing and visualising data to discover relationships and correlations, gain an overview of data distributions and unveil trends. Visual interfaces leverage the enormous bandwidth of the human visual system to support pattern discovery in large amounts of data. Our Knowminer search builds upon the well-known faceted search approach which is extended with interactive visualisations allowing users to analyse different aspects of the result set. Additionally, our system provides functionality for organising interesting search results into portfolios, and also supports social features for rating and boosting search results and for sharing and annotating portfolios.

Keywords—Search interface, faceted search, multi-visualisation analysis, shared result portfolios.

I. INTRODUCTION

The amount of information available on the internet, on our computers as well as on mass storage devices is growing every day. Hence, discovering the desired information is becoming a more and more complex task. Thus, the demand for powerful search tools, supporting efficient filtering and analysis of the result set, has increased significantly in recent times and became a major demand.

Marchionini [17] distinguishes two ways of searching depending on the task: information lookup and exploratory search. Information lookup can be satisfied quite easily by simply providing search results which is supported by standard web or enterprise search engines. Exploratory search however can be divided into learning and investigation tasks [10] and is characterised by taking a longer period of time involving much more interactions compared to simple information lookup tasks.

Our search interface enables both types of searching. A faceted search interface provides the possibility to efficiently narrow down the search result set. If users desire, they

can engage in further exploration and analysis using the visualisations.

Analysing the content becomes hard if the semantics of the analysed data is not available for machine processing. Tim Berners-Lee proposed the concept of the semantic web [29] where data is described by machine-readable semantic information. However, this approach has so far gained limited traction, partly because of the complexity of the associated technologies (such as RDF and SPARQL) and partly due to limited amount of high quality semantic data. While our system does not follow the semantic web approach it does employ lightweight semantics. We employ our Knowminer framework [16] which provides services for cluster analysis, classification and information extraction [14]. The automatically extracted information is used to annotate the content and is indexed along with the available meta data and the full text. Thus, by retrieving the meta data and the extracted information we i) provide a faceted view allowing users to refine their search queries in order to receive more precise results, and ii) provide datatype-specific visualisations supporting explorative analysis in the data.

In this paper we present a visual-supported faceted search interface as proposed by Kaiser et al. [13]. The backend of our search solution is based on Apache Lucene [1] search engine, and the Knowminer framework for automated document analysis which allows us to extract persons, places, organisations, topics and other entities, and to establish relationships between them. In the web-based frontend the extracted entities are grouped by type and displayed as facets in a tree which integrates bar charts to represent the document frequency of each entity. Furthermore, we provide several views visualising entities and documents in different contexts: i) a geo-visualisation shows the distribution of extracted geo-references, ii) a streamgraph displays trends and correlations between facets, iii) a graph visualisation enables the exploration of relationships between entities and documents, and iv) the information landscape gives an overview of the topical composition of the search result set. Besides providing means for explorative analysis the visualisations also support filtering operations on the data.

For organising and structuring search results we provide the portfolio feature, which allows users to assign important hits to topical collections. Each portfolio can be retrieved at a later time-point for in-depth analysis. Our system also provides social features like rating and boosting of search results as well as tagging and sharing of portfolios.

The unique feature set provided by our interface makes it suitable for a wide range of discovery and analytical tasks, such as in the fields of competitive intelligence, intellectual property or media analytics. Identifying competitors or partners working in the same area, discovering involved persons and relationships between them, or understanding temporal correlations between organisations and research topics are just a few examples.

This paper is structured as follows. Section II provides an overview of existing search engines as well as research done in this area. Section III gives a short overview on the backend implementation. Section IV describes the different components of the user interface including the search results, the facet view, the four visualisations and the portfolio view. Section V provides an overview of our explorative user study and discusses the results from this evaluation. Finally, the conclusion summarises our approach and the evaluation results, and discusses future work.

II. RELATED WORK

As the demand for powerful search engines has grown considerably, the number of available solutions supporting search and exploration of the data has increased over the last years. Many search solutions are domain specific and cannot be applied for a general purpose. Dessault System [7] provides meta data comparison, hierarchy navigation and 2D/3D views for searching reusable parts in the hardware industry. IBM's Content Analytics tool [11] provides several views including correlation and deviation views to analyse standard forms. Oracle's Endeca Information Discovery tool [23] allows analysing sales values by dynamically configuring the interface containing several standard visualisations. The IBM Data Explorer [12] is a search platform as well as an analytic tool for observing and analysing sales values using several chart visualisations. All these solutions provide visualisations which help users in investigating and analysing data in the specific domain.

Google Search Appliance [9] is a search product which can be applied for general purpose. Besides providing standard charts and a geo-visualisations, entity recognition, sentiment detection and question recognition are also supported. Another general applicable approach is Microsoft FAST [18] which provides simple charts e.g. pie or bar charts. Both mentioned general approaches provide simple visualisations with few interaction possibilities resulting in limited capability for explorative analysis of results. Cognito [5] provides several visualisations to explore the result set including a tag view, a graph view

as well as a geo-visualisation. Another enterprise search software has been proposed by Attivio [3]. In order to support visual analysis of the data Attivio connects Tableau [27] to their system [28]. Although this allows a detailed analysis of the result set, different view and programs are needed, which may lead to a loss of context.

Besides the above described search solution a lot of research has been performed recently in this area. Hearst [10] investigated visualisations in search interfaces and identified four different categories: visualising boolean syntaxes (e.g. venn diagrams), visualising query terms within retrieval results (e.g. tile bars), visualising relationships among words and documents (e.g. graphs or landscapes) and visualisation of text mining results (e.g. word tree or streamgraph). SemaVis [21] is a framework for searching and visualising semantic data in order to reveal hidden relations. SemaVis does not provide a classic search interface as it does not show a list of search results. Instead, the purpose is to display entities and relations between them in different visualisations. SemaVis employs adaptive visualisations depending on the user profile and past activities. Ding et al. [6] proposed a semantic web portal providing search functionality to display results in a facet view and in semantics-based visualisations (e.g. graph view, timeline and a geographic map) which are limited to RDF triples. RDF triples are used for conceptually describing web content and consist of subject, predicate and object. Foo et al. [8] provide a desktop search engine containing six different visualisations (tree, graph view, bubble chart, tile chart and cloud view) for analysing the results. However, the provided interaction possibilities for search refinement within the visualisations are limited. Arnaud et al. [2] provide a system for collaboratively organising, collecting and exploring sets of documents from different sources. In order to allow this, they use multi-scale visualisations providing automatic and interactive clustering. Clarkson et al. [4] provide a treemap-based visualisation to give an overview of existing meta data distribution. Another visual approach has been suggested by Nguyen et al. [22], where the solar system metaphor is used for visualising the result set. Locations, colours, movements and spatial distances are used for indicating relevances between a search query and the resulting web pages. Several interaction methods allow them to change the focus and provide insights. As the availability of sensory data is constantly growing Paparrizos et al. [24] provide a search interface for exploring sensor meta data including several visualisations (pie chart, bar chart and geo-visualisation). Kienreich et al. [15] proposed ApaLabs, a web-based visual-supported platform for retrieving and analysing news articles. ApaLabs provides several partially domain-specific visualizations including a geospatial visualisation, a tag cloud, and a parliament seat distribution visualisation of the Austrian National Council. Sabol et al. [26] provided a visual

analytics tool including an information landscape, a stream view as well as a cluster tree view in order to allow users to analyse topical-temporal relationships in large data sets. The method is supported by a highly scalable, incremental clustering approach to generate a hierarchical geometry for the visualisation [19].

III. BACKEND

In this Section we briefly describe the main server-side components. The basic components of the server are: the *data crawler and importer* for crawling and importing from different sources (e.g. databases, media wikis, file systems, web domains etc.), the *data converter* for converting the different file formats to a common format for text and meta data which can be consumed directly by the search engine, the *information extractor*, which enriches the textual content with semantic information like language, dates, geo-coordinates, named entities etc., the *classifier* which assigns documents to a number of pre-defined classes, the *clusterer* which computes topical clusters described by keywords from the underlying documents, the *information retrieval* component based on Lucene and extended by in-house features for faceted search, query term suggestions, spelling corrections and others, the *association retrieval component* which computes relationships between entities depending on their co-occurrence in the documents, and the *user management* for storing user specific data. Our search engine has been developed as a plug-in architecture to allow different server configurations depending on the demands. For example, plugins can be used for query expansion, ordering of facets to ensure that the most important entries are displayed first, or providing additional data required by the visualisations.

IV. SEARCH INTERFACE

In this Section we provide an overview of our search interface which is composed of the search result list, the facet view, the portfolio view and the visualisations.

The search interface is split up into several areas, see Fig. 1. The upper area provides a query input field and a drop down menu for index selection. The latter provides the possibility to search in different indices, each with different access rights. Just below the search area, the filter bar is displayed showing all selected facets, which are added to the original search query for filtering the results. On the left side, facet categories (e.g. persons) are shown as collapsible tree nodes containing the most frequently mentioned facets. The search result list is placed centrally, while the optional portfolio view is positioned on the right. Four buttons representing the visualisations are placed between the facet bar and the result list. Opening the visualisation moves down the search results and shows the selected visualisation in a tab view.

A. Search Results

After searching, the ten most relevant search results are displayed, further results can be accessed on paging. Initially these results are sorted according to their relevance but can also be sorted according to their average user rating or their modification date.

Each search result consists of the title, the modification date and a snippet. The snippet contains a small part of the content which may refer to the search query and helps in deciding whether the result is relevant for the selected query. Additionally several features are provided, e.g. boosting (raising relevance of result for current query), rating, searching for similar results and adding the result to a portfolio.

B. Facets

Facets are extracted meta data from the search results. Facets are grouped respectively to their type. Bar charts in the background of each facet as well as the corresponding number in brackets indicate how many results contain the specific facet. Each facet type has a different colour. The same colour coding is used within the visualisations to help users in identifying different types faster. Clicking on a facet refines the search and narrows down the number of results. Afterwards this facet is highlighted and added to the facet bar below the search input. Depending on the displayed visualisation, further interaction possibilities are provided. Additionally, search refinement can be done by filtering results by their rating.

Facet filters can be removed individually or all at once.

C. Visualisations

Visualisations provide additional information about the search results which are not conveyed by facets and the result list.

Geo-Visualisation: The geo-visualisation shows a map displaying the distribution of locations extracted from the search results, see Fig. 2. Locations are displayed through icons, whereby aggregation is performed where the icon density would be too high. The icon sizes indicate the number of aggregated locations. On zooming in, single locations may be revealed depending on the zoom level and their distance to nearby locations. Tool tips on aggregated location icons display the exact number of locations hidden by this icon, tool tips on single locations displays the three most relevant facets in the form of a bar chart. Search refinement within the geo-visualisation is performed by clicking on a location icon (single or aggregated) or by selecting an area within the map. In both cases the selected locations will be added to the facet bar.

Streamgraph: The streamgraph displays extracted information in the context of time, see Fig. 2. The uppermost area provides the possibility to add and remove facets from the visualisation, depending on users interests. The exact

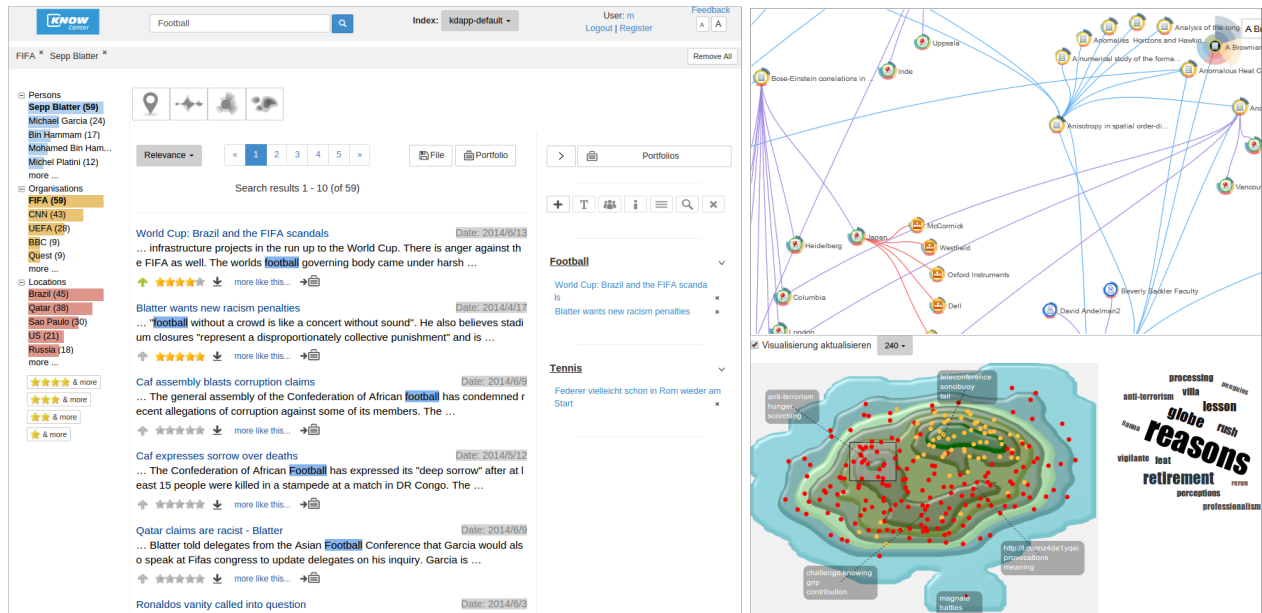


Figure 1: The web-based faceted search interface (on left), the graph view showing relations between nodes representing documents and entities of different types (up-right), and the information landscape providing a thematic overview of the result set (down-right).

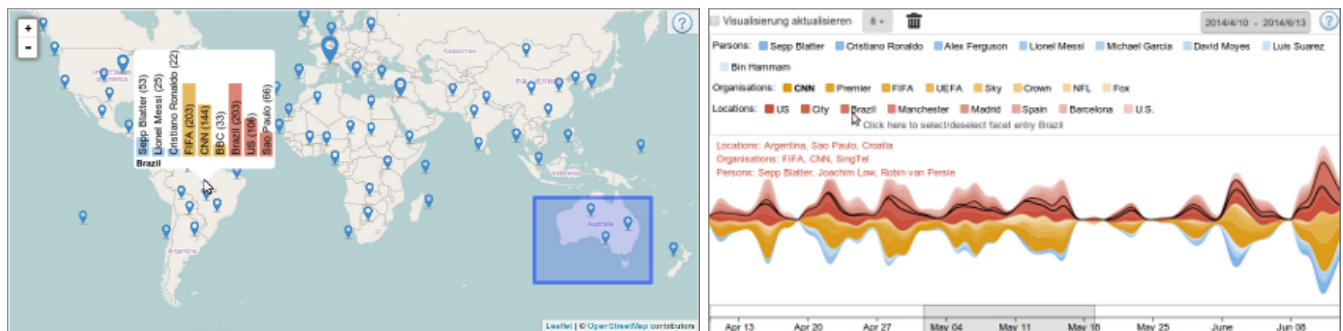


Figure 2: Geo-visualisation showing locations mentioned in the hits (left) and a streamgraph showing the number of documents mentioning each facets over time (right).

time span of the visualisation is displayed in the upper right. The legend below shows the displayed facets sorted by their type and by the number of documents mentioning them. Below the legend, the streamgraph shows the temporal information for the selected facets. The thickness of each stream corresponds to the number of documents mentioning the facet at the given point in time. Besides seeing the trends one can spot correlations in the temporal development of the facets. The time interval selection bar is displayed below the streams. The selected interval will be added to the facet bar and the results will be filtered accordingly.

Graph View: The graph view shows relationships between entities and documents, see Fig. 1. Relations between documents and entities (type "contained in") are extracted by named entity extraction methods. Relations between documents (type "similar to") are retrieved by similarity search (search by example). Relations between entities (type "associated with") are computed by co-occurrence analysis. Initially the graph is populated with the six most relevant search results, where additional documents or facets can be added through drag & drop from the result list or the facet tree, respectively. On mouse-over a node is represented

by a semantic blossom [25] allowing the user to expand relations along a path they are interested in: by clicking on the corresponding leaf only nodes of the corresponding type will be expanded. Additionally, users can restrict the set of expanded nodes by taking in consideration the context (i.e. a subset of the nodes currently visible on screen). In this way queries can be formulated simply by expanding and selecting nodes, for example given an organisation, show all related persons which are also connected to some of the locations shown on screen. Clicking on a facet node refines the search and adds the facet to the facet bar, clicking on a search result node opens the corresponding page or document in the browser.

Landscape: The information landscape visualisation [30] provides an overview of the topics covered by the search results, see Fig. 1. Results with similar content are positioned close together (in clusters), whereas results with different content are in different clusters. Clusters, represented by islands and hills, arise in the landscape where the density of highly similar results is large. Descriptive labels are generated for each cluster consisting of the most frequent keywords from the underlying documents. Tool tips on the red dots shows the title of the document, clicking on them opens the corresponding result. Clicking on a keyword within a label refines the search using this keyword. Inspection by selecting an area within the visualisation calculates the most important keywords for the selected results and displays them as a tag cloud. Also, mouse-over in the facet view highlights all results which contain the specific facet.

D. Portfolios

The portfolio view allows users to collect important results and assign them to different topical collection in order to access and analyse them later. Results can be added to portfolios by drag & drop. On portfolio selection, additional features are provided. Annotations can be entered by the users which will be used in a future version of our system for recommending results for portfolios. Furthermore users can share their portfolios with each others and maintain them collaboratively. Shared portfolios can either include full modification access or read only access. Additionally, we provide the possibility to visually analyse the documents contained within a portfolio.

V. EVALUATION

We performed a usability evaluation of our user interface to find out

- 1) whether the provided visualisations are helpful in analysing and refinement,
- 2) whether the provided interaction possibilities are sufficient,
- 3) whether users could apply the visualisations as easily as the facet view,

- 4) which visualisations and interaction possibilities are more demanding compared to others.

Seventeen volunteers participated in our evaluation, thirteen male and four female (average age was 32.8 years). The evaluation was performed using the thinking aloud protocol.

The tasks were split up into six task-groups: users first had to interact with the facet tree, then they performed tasks using the four different visualisations, and the final task addressed interactions with the portfolios. After each task-group was completed we collected information on the subjective task load using the the NASA Task Load Index [20] on a five-point scale (1 - lowest, 5 - highest). After completing all tasks users filled out an exit questionnaire.

A. Evaluation Results

The evaluation revealed that, considering the task completion rate, understanding and interacting with the geo-visualisation (59%) appears most difficult, followed by the streamgraph (82%) and the graph (84%), with almost equal completion rates, and the landscape (92%). The highest task completion rate was received using the facet tree (94%).

Table I shows the subjective task load results for the six task groups. Task load results confirm the task completion results. Interactions in the geo-visualisation were more demanding than using other visualisations. Table I further reveals that interactions in the graph and the streamgraph are about equally demanding whereas using the landscape was less demanding. The best (i.e. lowest, except for performance) values were reached on facet tree interactions due to the limited amount of interaction possibilities and the familiarity of using a tree view.

B. Personal Assessment

In this Section we present the results gathered by the questionnaires. According to the participants, the facet tree supports the refinement of a search query well, with the bar chart representation being helpful to recognise the relative importance of the facets (average rating > 4). Participants stated that they like the idea of portfolios which were appropriate for collecting documents, but provided many suggestions for additional features and design improvements. These include saving used search queries in order to reapply them later, providing additional search methods e.g. using portfolio annotations for search queries, retrieving suggestions for extending a portfolio, as well as the possibility to export portfolios to the local file storage. Furthermore users requested additional information on a portfolio, e.g. who shared a portfolio or the last modification date.

Considering the visualisations, users rated them as being helpful as they provide powerful additional possibilities to explore and filter the result set. A few users were concerned about the complexity level the visualisations add to the user interface, stating that the amount of interaction possibilities in the visualisations was overwhelming and that some

	mental demand	temporal demand	performance	effort	frustration
Facet tree tasks	1.4	1.5	4.3	1.5	1.1
Geo-visualisation tasks	2.6	2	3.3	2.7	2.4
Streamgraph tasks	2.4	2.1	3.4	2.2	1.9
Graph view tasks	2.4	1.8	3.7	2.3	1.9
Landscape tasks	2.2	1.9	3.9	2.1	1.8
Portfolio tasks	2.4	1.8	3.9	2.2	1.9

Table I: Average results on personal assessment concerning mental demand, temporal demand, performance, effort and frustration on a five point scale.

training is needed before using them. However as using the visualisations (and portfolios) is optional - they are hidden per default - we believe that our search UI design is not overloaded.

Asked about what they would use the visualisations for, participants stated that the geo-visualisation is useful for getting an overview of mentioned geographical references and refining the search by locations, while the streamgraph would be used for analysing changes over time and to identify correlations between different entities. The graph is primarily useful for inspecting meta data and extracted entities for search results, but can also be used to investigating relations between entities and finding similar results. The smallest advantage was identified for the information landscape: those participants willing to use it would just obtain an overview of the topical composition of the results set by inspecting the cluster labels.

VI. CONCLUSION

In this paper we presented a mature visually supported faceted search interface, backed by advanced text mining methods provided by the Knowminer framework, which allows users to gain deeper insights to the search results using interactive visualisations conveying temporal, geographical, topical and relationship information. We also provide a portfolio feature for organising and saving important results, and include social features for search results (rating and boosting) as well as for portfolios (tagging and sharing). We believe that our system provides a most comprehensive and unique combination of features, currently not supported by search engines available on the market nor in the recent research prototypes.

The evaluation results indicate that users were satisfied with the functionality of our search interface and they would use it for their particular purposes. Participants were able to analyse the result set using visualisations and discover patterns, however with comparably more difficulties using the geo-visualisation and a lower advantage seen for the information landscape. In general, the variety of interaction possibilities provided by the visualisations was perceived as overwhelming so that a training would be necessary to understand the functionality and the interplay between different views.

In the near future we will investigate ways to reduce the

complexity of the visualisations while retaining the most important features. We plan to extend the portfolio functions, for example by implementing automatic suggestions of portfolios for search result, providing advanced search using a portfolio as a query seed, and enabling users to compare different portfolios. We will also work on improving social features, such as tagging of search results, and plan to provide a query storage enabling saving and reapplying of complex queries and filters.

ACKNOWLEDGMENT

The Know-Center is funded within the Austrian COMET Program - Competence Centers for Excellent Technologies - of the Austrian Federal Ministry of Transport, Innovation and Technology, the Austrian Federal Ministry of Economy, Family and Youth and by the State of Styria. COMET is managed by the Austrian Research Promotion Agency (FFG).

REFERENCES

- [1] Apache Lucene. <http://lucene.apache.org/core/>.
- [2] B. Arnaud, G. Artignan, J. Cance, G. Delmas, M. Hascoët and N. Rodriguez. CoViz: cooperative visualization to facilitate sense making by groups of users. In Proceedings of the 15th International Conference on Information Visualisation (IV) 2011, pp. 255–260 .
- [3] Attivio. <http://www.attivio.com/>.
- [4] E. Clarkson, K. Desai and J. Foley. ResultMaps: Visualization for Search Interfaces. IEEE Transactions on Visualization and Computer Graphics. Los Alamitos, CA, USA, 2009.
- [5] Cognito. <http://www.expertsystem.com/cogito/>.
- [6] Y. Ding, Y. Sun, B. Chen, K. Borner, L. Ding, D. Wild, M. Wu, D. DiFranzo, A. G. Fuenzalida, D. Li, S. Milojevic, S. Chen, M. Sankaranarayanan and I. Toma. Semantic Web Portal: A Platform for Better Browsing and Visualizing Semantic Data. In Proceedings of the 2010 International Conference on Active Media Technology. Toronto, Canada, August 2010.
- [7] Exalead Onepart. <http://www.3ds.com/products-services/exalead/capabilities/digital-assets/exalead-onepart/>.
- [8] S. Foo and D. Hendry. Desktop search Engine Visualisation and Evaluation. Proceedings of the 10th International Conference on Asian Digital Libraries, ICADL 2007, Hanoi, Vietnam, 2007.

- [9] Google Search Appliance. <https://www.google.com/work/search/products/gsa.html>.
- [10] M. A. Hearst. *Modern Information Retrieval: The Concepts and Technology behind Search Engines* (2nd Edition).
- [11] IBM Content Analytics. <http://www-03.ibm.com/software/products/de/watson-content-analytics>.
- [12] IBM DataExplorer (Vivisimo). <http://www.ndm.net/datawarehouse/IBM/infosphere-dataexplorer-vivisimo>.
- [13] C. Kaiser, S. Softic, V. Sabol and M. Zechner. A User-Interaction Approach in an Information-Cockpit. In *Proceedings of the International Conference on Web Information Systems and Technologies, WEBIST 2013*, pp. 635–639.
- [14] R. Kern. *Efficient Data Structures for Knowledge Relationship Discovery*. Doctoral Dissertation, University of Technology Graz, 2012.
- [15] W. Kienreich, E. Lex and C. Seifert. APA Labs: An experimental web-based platform for the retrieval and analysis of news articles. *Proceedings of the First International Conference on the Applications of Digital Information and Web Technologies, ICADIWT 2008*, pp. 58–62.
- [16] W. Klieber, V. Sabol, M. Muhr, R. Kern, G. Öttl and M. Granitzer. Knowledge Discovery using the Knowminer Framework. *Proceedings of the IADIS International Conference Information Systems, IADIS*, pp. 307–314.
- [17] G. Marchionini. *Exploratory Search: From Finding to Understanding*. Communications of the ACM, 2006.
- [18] Microsoft FAST. <http://www.searchtechnologies.com/microsoft-search-services>.
- [19] M. Muhr, V. Sabol and M. Granitzer. Scalable Recursive Top-Down Hierarchical Clustering Approach with Implicit Model Selection for Textual Data Sets. In *Proceedings of the 2010 Workshop on Database and Expert System Applications, DEXA 2010*, pp. 15–19.
- [20] NASA Task Load Index. <http://humansystems.arc.nasa.gov/groups/tlx/downloads/TLXScale.pdf>.
- [21] K. Nazemi, M. Breyer, D. Burkhardt, C. Stab and J. Kohlhammer. SemaVis: A New Approach for Visualizing Semantic Information. *Towards the Internet of Services: The THESEUS Research Program. Cognitive Technologies 2014*, pp. 191–202.
- [22] T. N. Nguyen and J. Zhang. A Novel Visualization Model for Web Search Results. *IEEE Transactions on Visualization and Computer Graphics*, 2006.
- [23] Oracle Endeca Information Discovery. <http://www.oracle.com/us/solutions/business-analytics/business-intelligence/endeca/overview/index.html>.
- [24] I. Paparrizos, H. Jeung and K. Aberer. Advanced Search, Visualization and Tagging of Sensor Metadata. *Proceedings of the 2011 IEEE 27th International Conference on Data Engineering*.
- [25] M. Rauch, R. Wozelka, E. Veas and V. Sabol. Semantic Blossom Graph: A new Approach for Visual Graph Exploration. In *Proceedings of the 18th International Conference on Information Visualisation (IV) 2014*, pp. 234–240.
- [26] V. Sabol, W. Kienreich, M. Muhr, W. Klieber and M. Granitzer. Visual Knowledge Discovery in Dynamic Enterprise Text Repositories. In *Proceedings of the International Conference on Information Visualisation (IV) 2009*, pp. 361–368.
- [27] Tableau. <http://www.tableau.com/public/>.
- [28] C. Stolte, D. Tang and P. Hanrahan. Polaris: A System for Query, Analysis, and Visualization of Multidimensional Relational Databases. *IEEE Transactions on Visualization and Computer Graphics*, Volume 8, pp. 52–65, 2002.
- [29] The Semantic Web. <http://www.scientificamerican.com/article/the-semantic-web/>.
- [30] E. Ulbrich, E. Veas, S. Singh and V. Sabol. Reading through Graphics: Interactive Landscapes to Explore Dynamic Topic Spaces. In *Proceedings of the 17th International Conference on Human-Computer Interaction, Los Angeles, 2015*.