

# A visual information seeking system for Web search

Harald Reiterer, Gabriela Mußler, Thomas M. Mann

Computer and Information Science, University of Konstanz

## Abstract

In this paper we present the conception and the evaluation of a visual information seeking system for the Web. Our work has been motivated by the lack of good user interfaces assisting the user in searching the Web. The selected visualisations and why they have been chosen are explained in detail. A focus of this paper is on the evaluation of these visualisations as an add-on to the traditional result list presentation.

## 1 Introduction

Some of the main challenges of the Web are problems related to the user and his interaction with the retrieval system. There are basically two problems: *how to specify a query and how to interpret the answer provided by the system*. Surveys have shown that users have problems with the current paradigm of information retrieval systems for Web search simply presenting a long list of results (Zamir, Etzioni 1998). These long lists of results are not very intuitive for finding the most relevant documents in the result set. The empirical study by (Jansen et al. 2000) shows how users search the Web. The findings of this study are very interesting for designers of information retrieval systems for the Web. The main conclusion of this study is that Web search users seem to differ significantly from users of traditional IR systems (like online databases, CD-ROMs and online public access catalogues). As a field of future work the authors of the study address “[...] to support the continued research into new types of user interfaces, intelligent user interfaces, or the use of software agents to aid users in a much simplified and transparent manner” (Jansen et al. 2000, 226).

The above empirical findings motivated us to develop a new type of user interface for Web retrieval that supports the user in the information seeking process by providing selected visualisations in addition to the traditional result list. Systems combining the functionality of retrieval systems with the possibilities of information visualisation systems are called *visual information seeking systems*. An important aspect of visual information seeking systems is their possibility to visualise a great variety of document characteristics allowing the user to choose the most appropriate for his task.

This paper presents our main design ideas developing a visual information seeking system called INSYDER. Chapter 2 gives a brief explanation of the typical information seeking activities that could be supported by a visual information seeking system. In chapter 3 we discuss, with the focus on the visualisations, the new features of the INSYDER system. Chapter 4 presents our synchronised visualisation approach of Web search results and the results of a summative evaluation. Conclusions and an outlook are given in chapter 5.

## 2 Information Seeking on the WWW

One of the first steps when dealing with information seeking systems is to get an idea how to describe the information seeking process best. A good example for a high level task approach is the four phase framework for information seeking by (Shneiderman, Byrd, Croft 1997):

- Formulation: expressing the search
- Action: launching the search
- Review of results: reading messages and outcomes resulting from the search
- Refinement: formulating the next step

For the development of the INSYDER system we have chosen this framework, because from the user's point of view it covers all phases of the information seeking process in comprehensive way. Various other models describing the information seeking process can be found in (Hearst 1999).

## 3 New features of the INSYDER system

During the development of the INSYDER system it was not intended to develop new visual metaphors supporting the retrieval process. The main idea was to select existing visualisations for text documents and to combine them in a new way. Nowadays there are a lot of visualisations of search results in document retrieval systems available. Our selection of existing visualisations was based on the assumption to find expressive visualisations keeping in mind the user target group (business analysts), their typical tasks (to find business data on the Web), their technical environment (typical desktop PC and not a high end workstation for extraordinary graphic representations), and the type of data to be visualised (text documents). The primary challenge from our point of view was the intelligent combination of the selected visualisation supporting different views on the retrieved document set and the documents itself. The primary idea was to present additional information about the retrieved documents to the user in a way that is intuitive, fast to interpret and able to scale to large document sets.

Another important difference of our INSYDER system compared to existing retrieval systems for the Web is the comprehensive visual support of different steps of the information seeking process. The visual views used in INSYDER support the interaction of the user with the system during the formulation of the query (e.g. visualisation of related terms of the query terms with a graph), during the review of the search results (e.g. visualisation of different document attributes like date, size, relevance of the document set with a scatter plot or visualisation of the distribution of the relevance of the query terms inside a document with a TileBar), and during the refinement of the query (e.g. visualisation of new query terms based on a relevance feed-back inside the graph representing the query terms).

The retrieval aspects of the visual information seeking system INSYDER have not been in the primary research focus. Nevertheless the system offers some retrieval features that are not very common in today's Web search engines (cf. (Reiterer, Mußler, Mann 2000)).

It is for sure not new to combine visualisations and information retrieval aspects, but nowadays systems which do a dynamic search with a document attribute generation and the different visualisations of these attributes and document inherent data are new. Our approach aimed



mation about retrieved documents to the user in a way that is intuitive, fast to interpret and which is able to scale large document sets. There are two categories of additional information that could be visualised: visualisation of document attributes, and visualisation of interdocument similarities. The second category intends to reduce the multidimensional document space to a 2-D or 3-D space so that the users will be able to visually perceive the relationships between documents. Typical approaches are document networks, spring embeddings, clustering, self-organising maps. Our approach is restricted to the first category: the visualisation of document attributes. One important effect is the possibility to group documents that share similar attributes. We have used two different approaches depending on the additional information presented to the user:

- Predefined document attributes: E.g. title, URL, server type, size, document type, date, language, relevance. The primary visual structures to show the predefined documents attributes are the Scatterplot and the Result Table.
- Query terms` distribution: This shows how the retrieved documents related to each of the terms are used in the query. The primary visual structures to show the query terms` distribution are the Bargraph, the TileBar and the Stacked Column.


The visual information seeking system INSYDER is not a general-purpose system like traditional search engines (e.g. AltaVista). Its context of use is to support small and medium sized enterprises (SMEs) of specific application domains finding business information on the Web. So the findings of general empirical studies like the above mentioned are useful in principle but have to be completed with more specific requirements. Therefore at the beginning of the project a field study has been conducted, using a questionnaire that has been answered by 73 selected companies (SMEs) in Italy, France and Great Britain. The aim was to understand the context of use (ISO 9241 Part 11) following a human-centred design approach (ISO 13407). The following requirements are based on this field study. The typical users of the application domains of the INSYDER project are experts in one of the two business domains: building and construction or CAD software. They are typically no experts in information seeking using information retrieval systems. They know the Web and have some limited understanding of search engines. Based on the experiences of the field study different task scenarios using an information seeking system like INSYDER to find business information have been developed. Our field study shows that the formulation of the information needs is normally done in an unstructured text. The typical technical environments of the users are business PCs. The study has shown that the processing power, the RAM, and the size of the screen is limited. So we had to keep this in mind and therefore it was not possible to use sophisticated 3D visuals structures only available on high-end PCs or special workstations.


Therefore the visual mappings of web documents we have chosen are text in 1D (Bargraph, TileBar, Stacked Column) and text in 2D (Scatterplot). This final selection of the visual structures was based on the above suggestions of the field study, an extensive study of the state of the art in visualising text documents (Mann 1999) and the design goal to orientate our visual structures as far as possible on typical business graphics. The field study shows that all users have a good understanding of this kind of graphics and use them during their daily work (e.g. in spreadsheet programs). Similar conclusions, mainly based on an overview of the research done in the area of visualisation of search results in document retrieval systems, can be found in (Zamir 1998).


Another important design decision was to use a *synchronised multiple view* approach. It of-

fers the user the possibility to choose the most appropriate visualisation view for his current demand or individual preferences. Our approach has similarities with the idea of "Multiple Co-ordinated Views" with "Snap-Together Visualisation (STV)" (North et al. 1999), e.g. offering the user co-ordinated views for exploring information. Other examples of visual information seeking systems following a multiple view approach can be found in e.g. (Ahlberg et al. 1995).

But there are still some drawbacks in this multiple view approach: The user interface of the system becomes more complex and therefore could be harder to use. The user can choose an inappropriate visualisation for a specific situation. To intercept the possible drawbacks a number of guidelines have been considered. The visual structures have been adapted to each other in colour, orientation and the overall style. The visualisations are synchronised in a way that a selection in one representation of the search result set will be updated immediately in the other representations.

The default presentation of the search results is a so-called  *Result Table*. The main difference to the traditional result list of search engines is that all available attributes of each document are shown in different columns of the table. Each row shows one document. The user has the possibility to sort each document attribute in an increasing or decreasing order or to customise the table to his personal preferences (e.g. to show only the attributes he is interested in or to rearrange the order of the columns). Figure 1 shows an example of the Result Table for the query *visual data mining* and its results. On the left there are the different Spheres of Interest, while on the right the user sees the Result Table and an integrated browser which shows a preview of the actually selected document.

The use of the  *Scatterplot* was mainly inspired by the visual information seeking systems Envision (Nowell et al. 1996) and FilmFinder (Ahlberg et al. 1994). A scatter plot is a 2D visual structure showing two variables of each document. Each document is represented by a blue coloured dot. The X and Y dimensions encode two document attributes. There are three predefined scatter plots available, each with a fixed definition of the X and Y dimension (see Figure 3): date/relevance; server type/number of documents; relevance/server type. The user also has the possibility of defining his X and Y dimensions using the available attributes for each document. The Scatterplot offers the user an easy way to navigate through the document space on the set level to find interesting search results. In Figure 3 the user has chosen the dimension server type classification on the y-axis and document type classification on the x-axis. It can be seen that most documents of the query come from commercial sites. The documents are classified into 5 categories. A square-box with a numeric label (number of documents) represents a document group (e.g. belonging to the same category or having the same relevance). Moving the mouse over a dot or square-box launches a tool tip showing important document attributes like title, size, date, category, and an abstract. Tool tips are available in all visual structures. To support the exploration of the search results the user can select documents or groups of documents via mouse click and a double-click on a dot launches the selected document in the Web browser.

The use of the  *Barchart* was mainly inspired by the work of (Veerasingam et al. 1995). The original idea of Barcharts has been adapted in several ways. Firstly, to have the same way of displaying the documents like in the other views where document details are given (Result Table, TileBar, Stacked Column), a horizontal orientation has been chosen. Therefore the Barchart is rotated by 90 degrees: top down instead of from right to left. Secondly the impres-

sion of a document as an entity is emphasised by using Gestalt principles, without disturbing the keyword orientation too much. The colours used for the different keywords are the same as for TileBars and Stacked Column. Each Barchart represents one document and shows the distribution of the relevance for each keyword of the query and the total relevance for the document. Therefore it is easy to detect if a document deals with one or more of the different keywords of the query. The headings of each column can be sorted in an increasing or decreasing order. This function offers the user the possibility to view the distribution of the relevance of each keyword individually. Figure 4 shows the same document collection as in the Scatterplot view. The red dots at the beginning of each line symbolise that these documents have been selected (e.g. in the Scatterplot view). The visualisation shows that the first document in the view seems to be the most relevant one, as all three keywords that have been searched for appear in the document with high relevancy. As described above, also in this view documents can be selected or deselected.

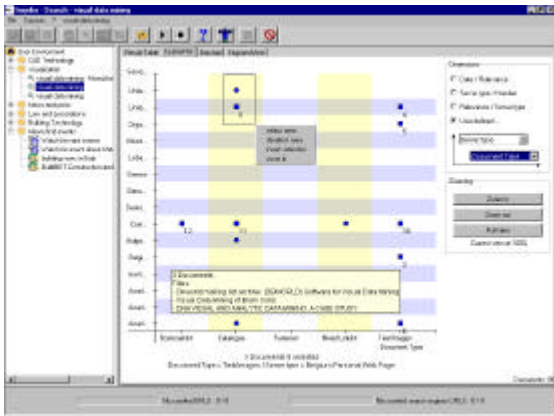


Figure 3: Scatterplot

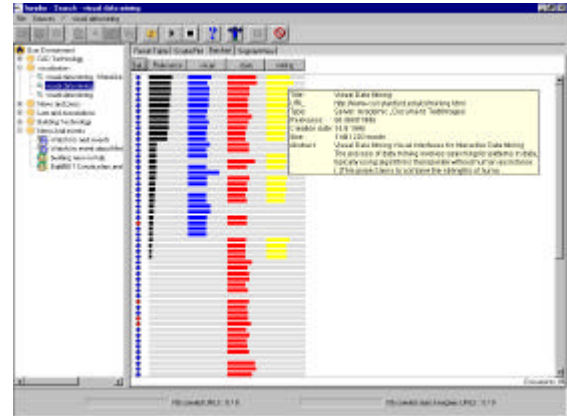



Figure 4: Barchart view

The visual structure  *TileBar* allows a deeper visual analysis on the document level, whereas the Scatterplots and Barchart are helpful on the document set level. The use of the *TileBar* was mainly inspired by the work of (Hearst 1995). Each document is represented by a rectangular bar, which is displayed next to the title of each document. The length of the rectangle indicates the length of the document. The bar is subdivided into rows that correspond to the keywords of the query. Each keyword is represented by a different colour. The bar is also subdivided into columns, where each column refers to a segment within the document. Keywords that overlap within the same segment are more likely to indicate a relevant document than keywords that are widely dispersed throughout the document. Thus, the user can quickly see if some subset of keywords overlaps in the same segment of the document. Figure 5 shows the display variant with 3 colour steps for the search with the three keywords "visual data mining". As shown by the blue, red and yellow tiles in the selected first document (title: Visual Data Mining) there are three segments with a co-occurrence of "visual" and "data" and "mining". If the user puts the mouse pointer over a segment, a tool tip occurs showing the text of this segment. A jump-into-segment feature for quick-jumps to the document-parts represented by the segments is also available. Figure 6 shows this jump-into-segment feature: a pop-up window occurs when clicking the right mouse button in the selected segment. The text of the segment is highlighted and put in the context of the whole document.

The *Stacked Column* visualisation is an adapted version of the Relevance Curve used in the

Result Table. Each segment of a document is represented as a vertical column. The height of each column corresponds to the number of times all keywords occur in that segment of text. For each keyword a different colour is used to show the contribution of each keyword to the overall occurrence of each column.

For evaluation purposes there is also a static  $\overline{\text{HTML}}List$  presentation in HTML-Format available in the INSYDER system (see section 4.3).

In all different views we have made extensive use of different *Interaction* techniques (e.g. direct manipulation, details-on-demand, zooming, direct selection) to give the user control over the mapping of data to visual form.

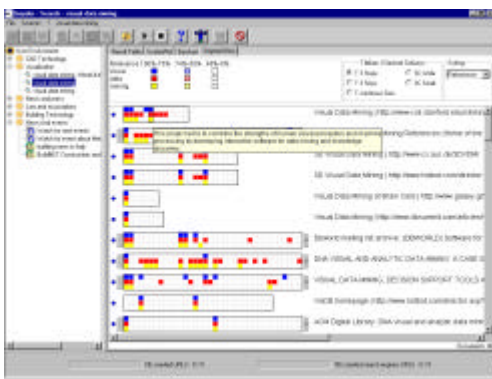


Figure 5: TileBar view

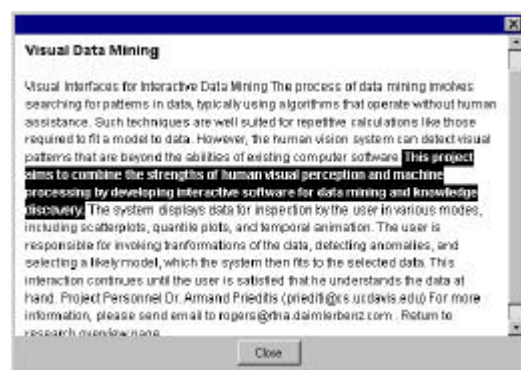


Figure 6: TileBar view - pop-up window

### 4.3 Evaluation of the Visualisations

The primary goal of the summative evaluation was to measure the added value of our visualisations in terms of effectiveness (accuracy and completeness with which users achieve task goals), efficiency (the task time users spent to achieve task goals), and subjective satisfaction (positive or negative attitudes toward the use of the visualisation) as dependent variables for re-viewing Web search results. Knowing advantages of the multiple view approach documented in user studies (North et al. 1999), we didn't intend to measure the effects of having Scatter-plot, Bargraph and TileBar/Stacked Column (also called SegmentView) *instead* of the List and Table. We wanted to see the added value of having these visualisations *in addition* to the Table and List.

From the factors influencing the design of a visual structure (Mann 1999) we decided to vary *target user group*, *type* and *number of data*, and *task* to be done. These have been determined as the independent variables. The fourth factor, the *technical environment*, was identical for all tests. The test setting covered all combinations of the different kinds of information seeking tasks (specific and extended fact finding), different kinds of users (beginners and experts), amount of results (30, 500), number of keywords of each query (1,3,8) and the chosen combinations of different visualisations (see Table 1).

#### 4.3.1 Procedure

A short entry questionnaire was used to record demographic data of each user. Then each user got a standardised system demo using a predefined ScreenCam recording presenting each



visualisation. After that each user had about 10 minutes to use the system and to ask questions if he had problems using it (learning period). The users were then asked to answer the 12 test task questions as quick as possible. All users processed the same 12 questions with the same keywords and number of hits in the same order, beginning with the 1 keyword / 30 hits specific fact finding task, followed by an extended fact finding 3 keywords / 500 hits question, then followed by a specific fact finding task with 8 keywords / 30 hits, and so on, always alternating specific fact finding and extended fact finding, with the number of keywords changing between every question in the order 1 – 3 – 8 – 1 – 3 – 8 ... . The presented visualisations to answer the questions have been different between the five groups (see Table 1). Based on this planning of the controlled experiment we could assure that the five combinations of visualisations have been distributed in an equal manner to all variables. During the tasks the users were requested to "think aloud" to enable the evaluation team to understand and record their current actions. The recording of data was done with written records taken by two persons. An experimenter moderated the test session so that in the case of problems this person could help. After accomplishing the test tasks the user had to answer a questionnaire of 30 questions regarding their subjective satisfaction and to suggest improvements of the system.

Question	Fact finding	collection size	# query terms	group 1	group 2	group 3	group 4	group 5
1	Specific	30	1	Table	Table	Table	Table	Table
2	Extended	500	3	Table	Table	Table	Table	Table
3	Specific	30	8	Table	Table	Table	Table	Table
4	Extended	500	1	Table	Table	Table	Table	Table
5	Specific	30	3	Table	Table	Table	Table	Table
6	Extended	500	8	Table	Table	Table	Table	Table
7	Specific	500	1	Table	Table	Table	Table	Table
8	Extended	30	3	Table	Table	Table	Table	Table
9	Specific	500	8	Table	Table	Table	Table	Table
10	Extended	30	1	Table	Table	Table	Table	Table
11	Specific	500	3	Table	Table	Table	Table	Table
12	Extended	30	8	Table	Table	Table	Table	Table

Table 1: Test combinations




Figure 7: Expected added value of visualisations expressed by users

### 4.3.2 Results

*Added values of the visualisations:* Figure 7 shows that in most test cases the users made use of the visualisations (using only the visualisation or using it in combination with the Result Table to answer the test questions). From this we conclude that the majority of the users expected an added value of the visualisations.

*Effectiveness:* The effectiveness of the visualisation is measured with the help of the degree of fulfilling the test tasks. E.g., if 8 out of the 12 tasks were solved, the effectiveness is 66,6% out of a maximum of 100%. As it can be seen in the left part of Figure 8 (Effectiveness) there was no significant advantage of using a specific visualisation combination. All visualisations performed nearly as good as the static list, which was used for reference purposes.

*Efficiency:* The efficiency of the visualisations has been defined as the effectiveness divided by the time the test persons needed to fulfil a test task. As no absolute minimum or best time exists for this test setting, the values derived are only comparable to each other. In Figure 8 (middle part) it can be seen that the Barchart combination performed second of all visualisation combinations. If we take into account that the Static List is something familiar to the user (well known from search engines), the Barchart has an outstanding role. Surprisingly it performs worst when looking at the effectiveness, but as the values are in a small interval, we do



not give too much strength to this effect. Taking the efficiency values, a first interpretation would be that training effects influence the use of the Barchart the least. Also the fact that the subjects expected a high added value from using a Scatterplot combination, but a low value in effectiveness and efficiency can be taken as a hint that training effects could have a high influence on the results. This will have to be evaluated in a next step. *User Satisfaction*: The user satisfaction is derived from the final questionnaire based on a Likert-scale (-2 to +2). Therefore positive and negative values occurred. For the user satisfaction an overall value has been calculated.

Figure 8 shows that this general impression of the visualisation was satisfying. This means that the majority of the test persons thought that none of the visualisations are dispensable. They also had the impression that the visualisations helped them to solve a task. The subjective impression of the Scatterplot was the worst.

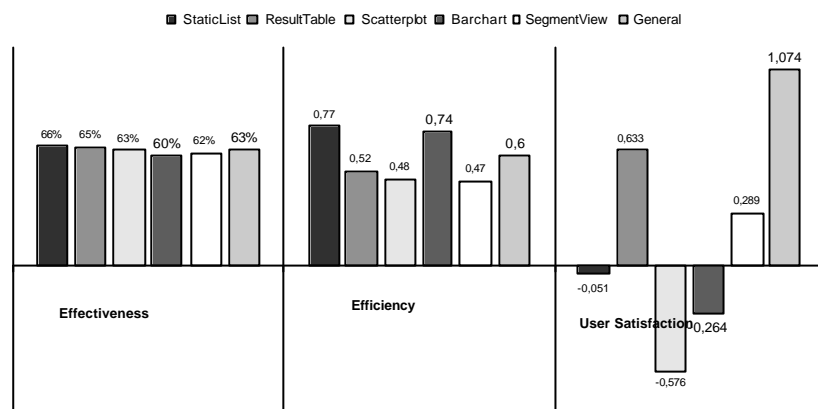


Figure 8: Overview of dependent variables

Users might have performed better, if they would have had more training time for the use of the Scatterplot and by performing better, it is likely that they have a more positive attitude towards a distinctive visualisation. Interestingly, most of the test persons were in a better mood after using INSYDER (positive mood before the test 92,5%, after the test 97,5%).

*Influence of target user group, type and number of data, type of task.* The numbers of documents, the numbers of keywords, the type of users, and the task type have shown to influence the efficiency of the visualisations.

## 5 Conclusion and Outlook

The results of the evaluation of our visual information seeking system for the Web have motivated us to go ahead. Our main design ideas for the development of a visual information seeking system for searching the Web have been successful. Most of the users make use of our synchronised multiple visual views and regarded them a nice enabling technology to find the most relevant documents in the search result. The evaluation results have shown that effectiveness and efficiency do not really increase when using visualisations, but the motivation and the subjective satisfaction do. We assume that more training time is needed to use the system effectively and efficiently.

Throughout the ideas presented above we are still working on the enhancement of the overall system. This includes the visualisations of the search results, developing specific filter functions supporting dynamic queries in combination with our visualisations, the visualisation algorithms and particularly the user interface of the whole application.

## References

- (Ahlberg et al. 1994) C. Ahlberg; B. Shneiderman. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. In: Proc. ACM CHI'94 pp. 313-317.
- (Ahlberg et al. 1995) C. Ahlberg; E. Wistrand. IVEE: An information visualization and exploration environment. In: Proc. IEEE Information Visualization 95, pp. 66-73.
- (Hearst 1995) M. A. Hearst. TileBars: Visualization of Term Distribution Information in Full Text Information Access. In: Proc. ACM CHI'95; 59-66, 1995.
- (Hearst 1999) M. A. Hearst. User interfaces and visualization. Modern Information Retrieval. R. Baeza-Yates and B. Ribeiro-Neto (eds.). Addison-Wesley (New York): 257-323, 1999.
- (Jansen et al. 2000) B. Jansen; A. Spink A.; T. Saracevic. Real life, real users, and real needs : a study and analysis of user queries on the web. In: Information Processing and Management 36, 2000, pp. 207-227
- (Koenemann, Belkin 1999) J. Koenemann and N. J. Belkin. A Case for Interaction: A Study of Interactive Information Retrieval Behavior and Effectiveness. CHI 96 - Electronic Proceedings. R. Bilger, S. Guest and M. J. Tauber (eds.). [http://www.uni-paderborn.de/StaffWeb/chi96/ElPub/WWW/chi96www/papers/Koenemann/jk1\\_txt.htm](http://www.uni-paderborn.de/StaffWeb/chi96/ElPub/WWW/chi96www/papers/Koenemann/jk1_txt.htm) [1999-11-11].
- (Mann 1999) T. M. Mann. Visualization of WWW-Search Results. In: Proc. DEXA'99 Workshops, 1999.
- (Mann, Reiterer 1999) T. M. Mann, H. Reiterer. Case Study: A Combined Visualization Approach for WWW-Search Results. IEEE Information Visualization Symposium. N. Gershon, J. Dill and G. Wills (eds.). 1999 Late Breaking Hot Topics Proc. Supplement to: G. Will, D. Keim (eds.): Proc. 1999 IEEE Symposium on Information Visualization (InfoVis'99). Conference: San Francisco, CA, USA, October 24-29, 1999. Los Alamitos, CA (IEEE Computer Soc. Press). San Francisco 1999: 59-62. 1999.
- (Mann, Reiterer 2000) T. M. Mann, H. Reiterer. Evaluation of different Visualizations of WWW Search Results. Proc. Eleventh International Workshops on Database and Expert Systems Applications (DEXA 2000). Conference: Greenwich, UK, September 4-8, 2000 (IEEE Computer Society).
- (Nielsen 1997) J. Nielsen. Search and You May Find. <http://www.useit.com/alertbox/9707b.html> [1999-03-18].
- (North et al. 1999) C. North; B. Shneiderman. Snap-Together Visualization: Coordinating Multiple Views to Explore. University of Maryland, technical report CS-TR-4020 June 1999.
- (Nowell et al. 1996) L. Nowell; R. France; D. Hix; L. Heath; E. Fox. Visualizing Search Results: Some Alternatives To Query-Document Similarity. In: Proc. ACM SIGIR '96. pp. 67-75.
- (Pollock, Hockley 1997) A. Pollock and A. Hockley. What's Wrong with Internet Searching. D-Lib Magazine, 1997, <http://www.dlib.org/dlib/march97/bt/03pollock.html> [1999-02-01].
- (Reiterer, Mußler, Mann et al. 2000) H. Reiterer, G. Mußler, T. M. Mann and S. Handschuh: INSYDER - An Information Assistant for Business Intelligence. Proceedings of the annual International ACM SIGIR Conference on Research and Development in Information Retrieval SIGIR '00, Athens 24-28 July 2000.
- (Shneiderman, Byrd, Croft 1997) B. Shneiderman, D. Byrd, and W. B. Croft. Clarifying Search: A User-Interface Framework for Text Searches. D-Lib Magazine, 1997, <http://www.dlib.org/dlib/january97/retrieval/01shneiderman.html> [1999-08-17].
- (Veerasamy et al. 1995) A. Veerasamy; S. B. Navathe. Querying, Navigating and Visualizing a Digital Library Catalog. In: Proc. DL'95. <http://www.csd.tamu.edu/DL95/papers/veerasamy/veerasamy.html> [1999-03-24]
- (Voorhees, Harman 1998) E. M. Voorhees and D. K. Harman (eds.): NIST Special Publication 500-242: The Seventh Text Retrieval Conference (TREC-7) Gaithersburg, Maryland (Government Printing Office (GPO)) 1998. [http://trec.nist.gov/pubs/trec7/t7\\_proceedings.html](http://trec.nist.gov/pubs/trec7/t7_proceedings.html) [1999-12-20].
- (Zamir, Etzioni 1998) O. Zamir and O. Etzioni. Web Document Clustering: A Feasibility Demonstration. SIGIR 1998. <http://zhadum.cs.washington.edu/zamir/sigir98.ps> [1999-03-23].
- (Zamir 1998) O. Zamir. Visualization of Search Results in Document Retrieval Systems. General Examination's Paper, University of Washington, <http://www.cs.washington.edu/homes/zamir/papers/gen.doc> [2000-09-13]
- (Zizi; Beaudouin-Lafon 1994) M. Zizi; M. Beaudouin-Lafon. Accessing Hyperdocuments through Interactive Dynamic Maps. Conference on Hypertext and Hypermedia Proceedings of the 1994 ACM European conference on Hypermedia technology. 126-134.