Visual Information Retrieval for the WWW

Harald Reiterer, Thomas M. Mann, Gabriela Mußler

Computer and Information Science, University of Konstanz, Germany

Abstract

In this paper we present the conception and the evaluation of a visual information retrieval 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 the reasons why they have been chosen are explained in detail. An evaluation of these visualisations as an add-on to the traditional result list is presented.

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 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 retrieval systems*. An important aspect of visual information retrieval 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[1]. In chapter 2 we discuss, with the focus on the visualisations, the new features of the system. Chapter 3 presents our synchronised visualisation approach of Web search results and the results of a summative evaluation. Conclusions and an outlook are given in chapter 4.

2 INSYDER - a Visual Information Retrieval 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 (Hearst 1999). Our selection of existing visualisations was based on the assumption to find expressive visualisations keeping in mind the target user group (business analysts), their typical tasks (to find business data on the Web), their technical environment (typical desktop PC), 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, or relevance 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 feedback 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 (Reiterer et al. 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 docu-

[1] INSYDER (INternet SYstème DE Recherche) was funded by a grant from the European Union, ESPRIT project number 29232.

ment inherent data are new. Our approach aimed at getting the highest added-value for the user combining components like dynamic search, visualisation of the query and different visualisations of the results and information retrieval techniques (e.g. query expansion, relevance feedback).

3 Visualisations supporting the Information Retrieval Process

3.1 Visual Query Formulation

From the literature it is well known that users have problems formulating their information need (Pollock, Hockley 1997). This led to the demand of methods to overcome the problem of lacking knowledge to formulate queries. The idea of the visual query formulation is to help users to specify their information need more precisely using query expansion techniques and visualisation. The query expansion is implemented using a knowledge base, which is built upon terms and related concepts. Users can benefit from using it in two ways: either by changing their query terms leading to a more precise result set or by expanding their original query with additional terms from the knowledge base, which will result in a broadened result set, which could be much more satisfying, too. And as a side effect using terms from the knowledge base, spelling mistakes can be minimised, too. We propose a visual query, which will show the user related terms for his query (Figure 1), taking into account other successful solutions and ideas from automatic query expansion and query visualisation, e.g. (Voorhees, Harmann 1998), (Zizi, Beaudouin-Lafon 1994). As an intuitive way to express the relation of terms, we propose to use a graph for their visualisations. The entry point for the visualisations is the query entered by the user (e.g. *WWW visualisation*). The original and the resulting related terms are then presented in a graph and tree view. The graph view represents the terms with their "near" related terms. E.g. *WWW* is expanded to *Web, Internet, Media,* etc.. (see Figure 1). The tree view resembles the overall term space and therefor contains all terms related to a distinct term. These distinction is made to keep the graph view as easy as possible to survey. The broader related terms are displayed in the graph view using a hypertext metaphor for navigation: Clicking on a term (e.g. *Internet*) will show all other related terms for this distinctive term, depending on the number of all the related terms (e.g. *Hypertext, Usenet* etc.), it may come up as a circle with the expanded term as a centre (Figure 2). The two views are synchronised, which means that a term selection in the graph view will select the equivalent in the tree view and vice versa. If the graph view does not contain the selected term from the tree view, a new graph is created. Following the hypertext metaphor the visual query provides also a history function, so that the user can keep track of different graphs.

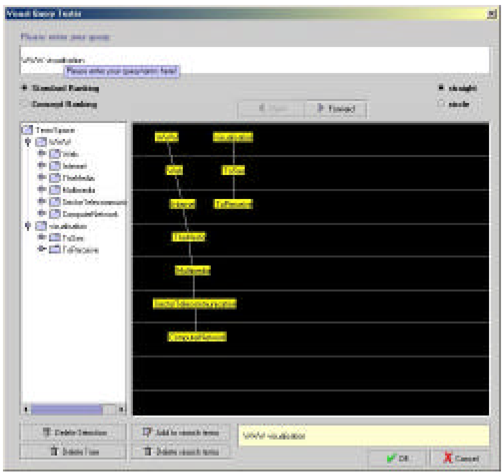
Figure 1: Visual query

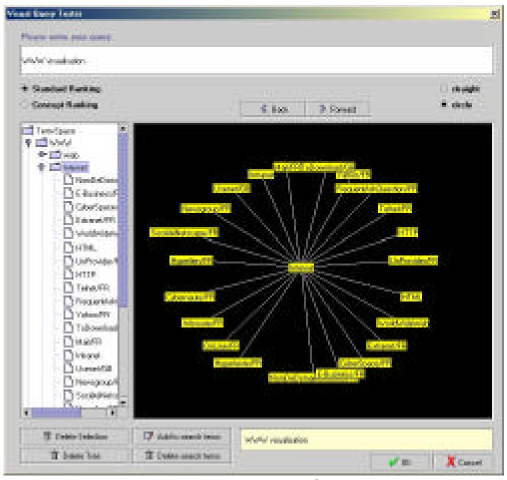
Figure 2: Examination of a term

The retrieval system uses two approaches to rank documents. The standard ranking takes all terms into account for its ranking, while the concept ranking takes the concepts from the knowledge base. These concepts can be weighted by the user to express his information need more precisely. In the visual query we therefore have also two ways foreseen. The standard follows the description above. The concept ranking (by selecting the appropriate radio button), will expand the graph view in two ways. Using the context menu the user can select the weight of each concept. A '+' sign is used as a marker for the importance, which will be shown above that term. According to the type of ranking terms are added to the search terms by using the appropriate button. The search terms are displayed in an own highlighted text entry field next to the button. Deleting the original search terms there (e.g. to take more precise

terms), the user can see them still in the upper entry field.
If the user uses the relevance feedback option of the system, the process is basically the same. The difference is that instead of the user the system provides the entry point terms for the graph and term space.

3.2 Visualisation of Search Results

The main idea behind our visual information retrieval approach is to present additional information about retrieved documents to the user in a way that is intuitive, fast to interpret and which is able to scale large document sets. One important feature 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 (a similar idea could be find in Ahlberg et al. 1994) and the Result Table (Figure 6).
- 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 (a similar idea could be find in Veerasamy et al. 1995), the TileBar (Hearst 1995) and the Stacked Column.

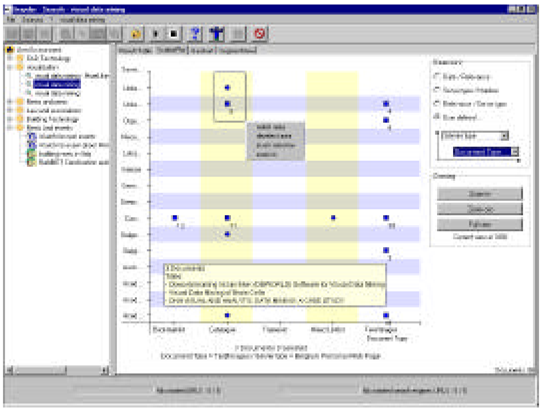

Figure 3: Scatterplot

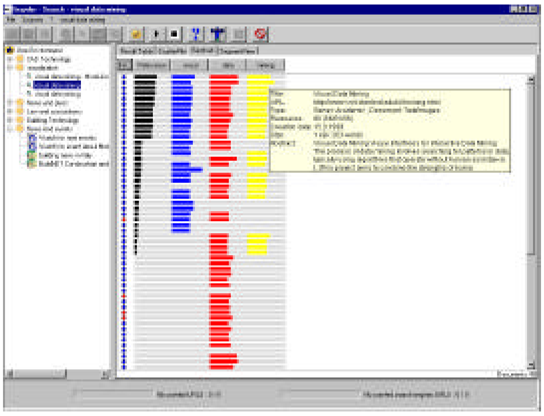

Figure 4: Barchart

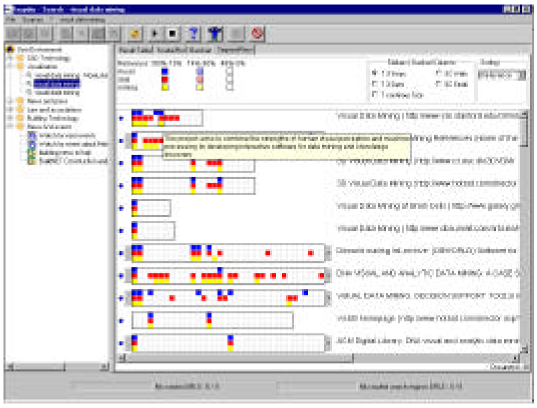

Figure 5: TileBars

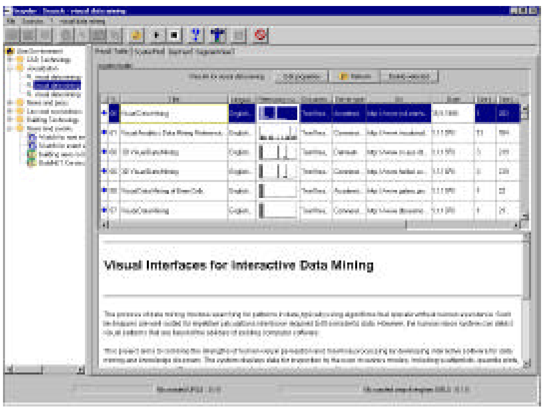

Figure 6: Result table

The visual mappings of web documents we have chosen are text in 1D: Bargraph (Figure 4), TileBars (Figure 5), Stacked Column and text in 2D: Scatterplot (Figure 3). This final selection of the visual structures was based on a field study, an extensive study of the state of the art in visualising text documents 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 offers 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 Coordinated Views" with "Snap-Together Visualisation (STV)" (North et al. 1999), e.g. offering the user coordinated views for exploring information.

3.3 Evaluation

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 reviewing 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 Scatterplot, 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, Reiterer 2000) 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. *Technical environment* and *training* 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.

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 get familiar with 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. During the tasks the users were requested to "think aloud" to enable the evaluation team to understand and record their current actions. After accomplishing the test tasks the users had to answer a questionnaire of 30 questions regarding their subjective satisfaction and to suggest improvements of the system.

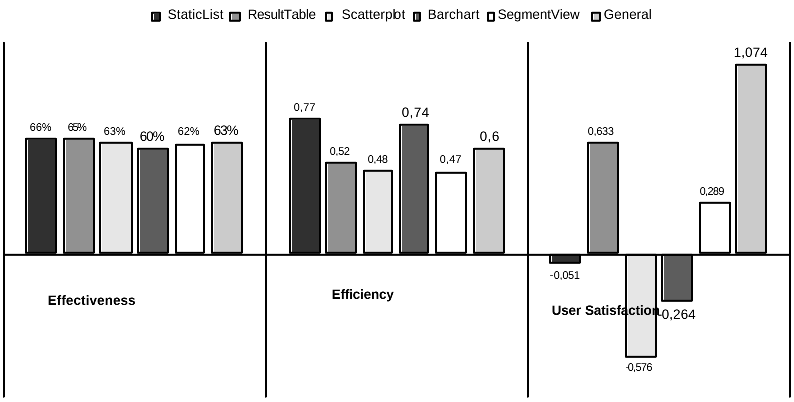

Figure 7: Evaluation results

Added values of the visualisations: 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 7 (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 7 (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. Also the fact that the subjects often used the Scatterplot combination, and therefore probably expected a high added value from using it, but had in realty 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 summarizing a number of questions. Figure 7 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. 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.

4 Conclusion and Outlook

The results of the evaluation of our visual information retrieval system for the Web have motivated us to go ahead. Our main design ideas for the development of a visual information retrieval 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.

(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.

(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). Con-ference: Greenwich, UK, September 4-8, 2000 (IEEE Computer Society).

(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.

(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 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 Con-ference on Research and Development in Information Retrieval SIGIR '00, Athens 24-28 July 2000.

(Veerasamy et al. 1995) A. Veerasamy; S. B. Navathe. Querying, Navigating and Visualizing a Digital Library Catalog. In: Proc. DL'95. http://www.csdl.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 [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 confer-ence on Hypermedia technology. 126-134.