Adaptive Search Suggestions for Digital Libraries

Sascha Kriewel and Norbert Fuhr

University of Duisburg-Essen

Abstract. In this paper, an adaptive tool for providing suggestions during the information search process is presented. The tool uses case-based reasoning techniques to find the most useful suggestions for a given situation by comparing them to a case base of previous situations and adapting the solution. The tool can learn from user participation.
A small, preliminary evaluation showed a high acceptance of the tool, even if improvements are still needed.

1 Introduction

A common and well known problem in the design of information retrieval (IR) and digital library systems is how to support end users in finding good strategies for satisfying their information need. Despite many advances in making information search technology available to the larger public instead of just search professionals, the effective use of these information retrieval technologies remains a challenge [8, 11, 19].

While all digital library and information retrieval systems provide low-level search actions, in [4] Bates identifies three higher levels of abstraction for categorizing search functionalities: tactics, stratagems, and strategies. There has been extensive work on supporting users in executing moves, tactics, and even stratagems, but there is a definite lack of support on the highest level [9, 10, 14], i.e. in helping searchers to choose the most appropriate or useful action for their specific situation to form an overall search strategy or plan.

In [17] we suggest that an ideal search system would provide useful strategies for completing specific tasks – either upon user request or pro-actively – and help users improve their search experience by raising awareness of the strategic aspect of searching. In this way, users would acquire procedural knowledge to create better search strategies for themselves.

In this work, a new tool is introduced that suggests appropriate search tactics and stratagems based on the current situation of the user. The tool, which has been implemented as part of the Daffodil framework [13], uses case-based reasoning (CBR) techniques to find and rank suggestions according to the similarity of the user's situation to previous situations. It also learns which of the several presented suggestions were successfully employed by users and which not.

2 Providing Strategic Help

For the purpose of this paper, we will follow the definition of a *search strategy* used in [4]. A strategy is seen as a complete search plan encompassing possibly many tactics and stratagems used in the process of an information search. To form such a search strategy, users need to be able to select the appropriate stratagems or tactics supported by the information system. It should be noted however, that searching is an opportunistic process and search goals can be shifting throughout the task. It is often more fruitful for a searcher to follow promising opportunities that arise from previous results, instead of sticking to a straight path towards the perfect result set. In [3] Bates describes the *berry picking* model of information seeking.

From the view of the user recognizing these strategic opportunities remains a problem, as does using the strategic options available to maneuver out of a perceived dead-end during a search. In fact, users rarely exploit the advanced capabilities and features of modern search systems, even if these would improve their search. They might not be aware of their existence, might not understand them, or don't know in which situations they could be effectively employed. Search systems providing strategic assistance could improve search effectiveness by suggesting the use of these advanced features or options automatically [15].

In [10] Brajnik et al. describe a strategic help system based on collaborative coaching, which tries to assist users by providing them with suggestions and hints during so-called critical or enhanceable situations. The system uses a hand-crafted knowledge base of 94 production rules to provide suggestions based on the tactics and stratagems proposed in [2, 4]. The strategic help module was integrated into FIRE, a user-interface to a Boolean IR system. Only six people participated in the user evaluation, but the results showed promise for the usefulness of strategic suggestions.

Belkin et al. describe information retrieval in terms of information-seeking behaviors or strategies (ISSs) [5, 7]. They present a characterization of such behaviors using a small set of dimensions. In the course of a single information-seeking episode, users will engage in several such ISSs, moving from one to the next. In a feature-rich search system like DAFFODIL[13] many of these ISSs correspond to specific supported stratagems or available tools, and the movement of the user between the various tools is similar to the movement from one ISS to another. Belkin et al. suggest a mixed-initiative system supporting specific procedures during information seeking episodes, which change and branch as the user interacts with the system. In [6] the MERIT system based on these concepts is presented. This system uses scripts derived by case-based reasoning techniques to guide users through an information-seeking episode.

Case-based reasoning emerged in the early 1980s as an Artificial Intelligence (AI) approach to solving a new problem based on the known solutions of similar problems. A stored case is a previously experienced problem situation which has been captured and retained together with its solution. New cases can be "solved" by comparing them to the case base of previous situations, finding the most similar ones, adapting them to the current situation, and then re-using them. Learning in CBR occurs as a by-product of problem solving. Whenever a specific case is successfully solved, this experience is retained in order to solve similar problems in the future (and accordingly for failures) [1].

3 The ASDL Suggestion Tool

To help users in their particular search situation with useful strategic search advice, the ASDL (Adaptive Support for Digital Libraries) Suggestion Tool was developed and integrated into Daffodil, an existing search system developed at the University of Duisburg-Essen [18]. Daffodil (pictured in Fig. 1) offers a rich set of tools with a correspondingly large number of possible user actions from which to draw suggestions. It also provides facilities for easily logging and gathering information about the search progress of users. Daffodil consists of an agent-based back-end and a graphical user client, which presents search, browse and extraction services in the form of desktop tools [12].

The ASDL module is composed of three main components:

- the *Observing Agent*, which collects information about the current search activities of the user and the results sent by the information sources;
- the *Reasoning Agent*, which retrieves appropriate suggestions fitting the user's current situation, and ranks them according to similarity;
- the actual *Suggestion Tool*, which adapts and presents the search suggestions within the Daffodil user interface, and allows for automatic execution and user judgments on suggestions.

3.1 Observing the User

For the purpose of providing search suggestions each completed user query with its corresponding set of result documents is considered as a single situation. The description for the situation is gathered from the original query of the user, the results sent by the different information sources, and additional information extracted by the Observing Agent. Thus a complete description of a search situation within the ASDL module contains

- the search terms, search fields, Boolean operators, and information sources used;
- the number of results returned and system response time;
- a list of the most frequent terms, authors, journals, or conferences extracted from the results.

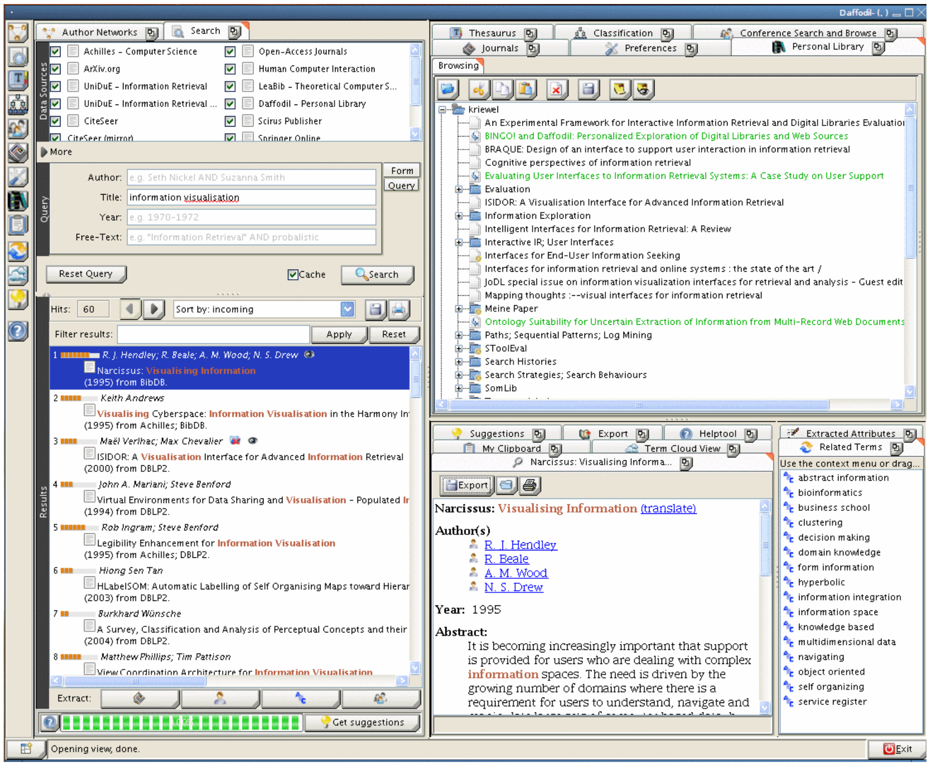

☑ Achilles – Computer Science
☑ ArXiv.org
☑ UniDuE – Information Retrieval
☑ UniDuE – Information Retrieval ...
☑ CiteSeer
☑ CiteSeer (mirror)
☑ Open-Access Journals
☑ Human Computer Interaction
☑ LeaBib – Theoretical Computer S...
☑ Daffodil – Personal Library
☑ Scirus Publisher
☑ Springer Online

Author: e.g. Seth Nickel AND Suzanna Smith
Title: information visualisation
Year: e.g. 1970-1972
Free-Text: e.g. "Information Retrieval" AND probalistic

☑ Cache

Hits: 60

Sort by: incoming

Filter results:

1 R. J. Hendley; R. Beale; A. M. Wood; N. S. Drew
Narcissus: Visualising Information
(1995) from BibDB.

2 Keith Andrews
Visualising Cyberspace: Information Visualisation in the Harmony In
(1995) from Achilles; BibDB.

3 Maël Verlhac; Max Chevalier
ISIDOR: A Visualisation Interface for Advanced Information Retrieval
(2000) from DBLP2.

4 John A. Mariani; Steve Benford
Virtual Environments for Data Sharing and Visualisation – Populated Ir
(1994) from DBLP2.

5 Rob Ingram; Steve Benford
Legibility Enhancement for Information Visualisation
(1995) from Achilles; DBLP2.

6 Hiong Sen Tan
HLabelSOM: Automatic Labelling of Self Organising Maps toward Hierar
(2003) from DBLP2.

7 Burkhard Wünsche
A Survey, Classification and Analysis of Perceptual Concepts and their
(2004) from DBLP2.

8 Matthew Phillips; Tim Pattison
View Coordination Architecture for Information Visualisation

100%

Opening view, done.

kriewel
- An Experimental Framework for Interactive Information Retrieval and Digital Libraries Evaluatio
- BINGO! and Daffodil: Personalized Exploration of Digital Libraries and Web Sources
- BRAQUE: Design of an interface to support user interaction in information retrieval
- Cognitive perspectives of information retrieval
- Evaluating User Interfaces to Information Retrieval Systems: A Case Study on User Support
- Evaluation
- ISIDOR: A Visualisation Interface for Advanced Information Retrieval
- Information Exploration
- Intelligent Interfaces for Information Retrieval: A Review
- Interactive IR; User Interfaces
- Interfaces for End-User Information Seeking
- Interfaces for information retrieval and online systems : the state of the art /
- JoDL special issue on information visualization interfaces for retrieval and analysis – Guest edit
- Mapping thoughts :--visual interfaces for information retrieval
- Meine Paper
- Ontology Suitability for Uncertain Extraction of Information from Multi-Record Web Documents
- Paths; Sequential Patterns; Log Mining
- SToolEval
- Search Histories
- Search Strategies; Search Behaviours
- SomLib

Narcissus: Visualising Informa...

Narcissus: Visualising Information (translate)

Author(s)
R. J. Hendley
R. Beale
A. M. Wood
N. S. Drew

Year: 1995

Abstract:
It is becoming increasingly important that support is provided for users who are dealing with complex information spaces. The need is driven by the growing number of domains where there is a requirement for users to understand, navigate and

Use the context menu or drag...
- abstract information
- bioinformatics
- business school
- clustering
- decision making
- domain knowledge
- form information
- hyperbolic
- information integration
- information space
- knowledge based
- multidimensional data
- navigating
- object oriented
- self organizing
- service register

Fig. 1. The DAFFODIL desktop: Search Tool (left), Personal Library (top right), Detail View and Related Terms (bottom right)

3.2 Finding Suggestions

The ASDL module uses case-based reasoning techniques to find situationally appropriate suggestions, where a search situation is considered as a case description and suggestions are possible solutions. Users can judge suggestions as useful for their current situation, thereby adding to the case base. The initial case base was build by creating one or two iconic cases for each suggestion.

For a given situation, the reasoning component tries to find the most similar cases from the database and returns the solutions, i.e. suggestions, ranked according to the descreasing similarity of the corresponding case to the current situation. If the same solution is suggested more than once, only the highest similarity is currently considered.

The similarity $sim_T(a, c)$ of two situations a and c is computed as the *weighted mean* (using weights w_k) of the individual similarity scores sim_k between the various aspects of a situation (see 3.1 above).

$$sim_T(a_i, c) := \frac{\sum_{k=0}^{N} w_k \cdot sim_k(a_{ik}, c_k)}{\sum_{k=0}^{N} w_k} \qquad (1)$$

For determining the similarity between two vectors of term weights (e.g. the extracted terms or authors from results), the normalized inner product is used.

3.3 Presenting the Help

When the results of a search are presented to the user, an unobtrusive button is shown below the list of results (visible in Fig. 1 to the right of the search progressbar). Upon clicking this button, the search suggestions are presented in form of a ranked list to the right of the result list (see Fig. 2). Each suggestion consists of a short descriptive title, a longer explanation and a small colored bar representing the estimated applicability of the suggestion to the current situation (i.e. the similarity score computed by the reasoning component).

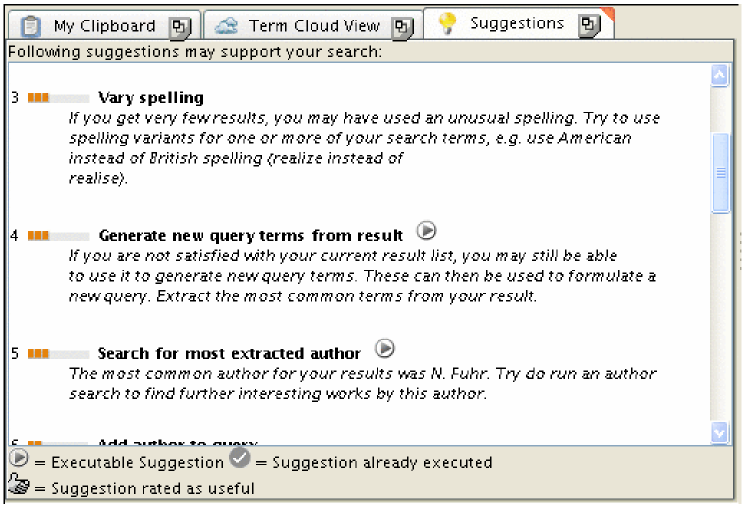

Following suggestions may support your search:

3 **Vary spelling**
If you get very few results, you may have used an unusual spelling. Try to use spelling variants for one or more of your search terms, e.g. use American instead of British spelling (realize instead of realise).

4 **Generate new query terms from result**
If you are not satisfied with your current result list, you may still be able to use it to generate new query terms. These can then be used to formulate a new query. Extract the most common terms from your result.

5 **Search for most extracted author**
The most common author for your results was N. Fuhr. Try do run an author search to find further interesting works by this author.

Fig. 2. The Suggestion Tool

Where possible, suggestions are adapted to the current search. E.g., a suggestion to perform an author search will suggested the most frequently extracted author name from the current results (if the query leading to the result was already an author search for that author, the second most frequently extracted author will be substituted).

A user can automatically execute most of the suggestions by double clicking on them, or by using the context menu. The list of suggestions remains visible, and additional suggestions can be tried. If a suggestion proves to be useful for furthering the search, the user can easily judge it as appropriate by using the context menu. Small icons are used to mark suggestions that can be automatically executed, that have already been executed in the current situation, or that have been judged as useful by the user.

3.4 Suggestions

A total number of sixteen suggestions were implemented for a first test, ranging from terminological hints (*vary spelling*) to suggestions for using different tools from the DAFFODIL toolbox (*show co-author network*). Wherever possible, the suggestions were made executable.

1. *Browse conference proceedings* using the most common conference from the current result set.
2. *Browse journals* using the most common journal from the current result set.
3. *Browse a classification* to find better terms to describe the information need.
4. Use the Thesaurus Tool to *Find a narrower term.*
5. Use the Thesaurus Tool to *Find a more general term.*
6. Use the Thesaurus Tool to *Find a related term.*
7. Use the Network Tool to generate a visual *co-author graph* for the most common author from the current result set.
8. *Perform an author search* using the most common author from the current result set (who was not already part of the current query).
9. *Restrict the query* by adding the most common *author* from the current result set (who was not already part of the current query).
10. *Restrict the query* by adding the most common *term* from the current result set (that was not already part of the current query).
11. *Restrict the query* by showing only the *most recent publications.*
12. *Restrict the query* by using the *phrase* operator for phrases.
13. *Broaden the query* by removing a search term from a conjunction.
14. *Broaden the query* by replacing the implicit conjunction with a disjunction.
15. *Vary spelling*, e.g. to compensate for differences in spelling between American and British English.
16. *Extract new query terms* from the current result documents and show them in form of a weighted list (also called *tag cloud*).

4 Evaluation

A pilot study was conducted to evaluate the ASDL prototype. It was a lightweight study with two main goals: gaining a first understanding of user acceptance of strategic suggestions, and evaluating the appropriateness of ranking suggestions according to previous users' situations and actions.

In preparation for the experiment the system was trained (using a search topic unrelated to that of the evaluation task), so that the case base contained at least one situation for each suggestion where it would be an appropriate advice. Additionally, no cut-off value for the similarity value was specified, so that all suggestions that could be applied to a situation were presented each time, even if their score was low.

A total of twelve participants were asked to perform a comprehensive search task with the help of the Suggestion Tool (ST): five graduate and seven undergraduate students of computer science or communication and media science from the University of Duisburg-Essen. They worked with a simplified version of the Daffodil system containing a sub-selection of tools, and performed their searches over several collections of computer science articles. The ST's purpose was described to them, and they were given an introduction to the Daffodil system where necessary.

The search task itself was only loosely defined, but comprehensive enough that even experienced searchers were expected to issue several queries and use a number of different tools from Daffodil's toolbox to gather a satisfying number of results. The participants of the evaluation were asked to "collect articles and additional information such as central authors, search terms, conferences or journal issues to prepare a course paper on methods for visualization of result lists in information retrieval systems". This search task was chosen to provide the searchers with a simulated work task and was only slightly modified from a real work task.

Each searcher was given between 40 and 50 minutes of time to work on the task, during which they were asked to use the suggestion tool for help and for generating new search ideas. During this time, one experimenter was available to give technical help and intervene on system errors. No strategic or terminological help was provided by the experimenter, and the searchers were referred to the Suggestion Tool instead.

Each session was logged using the logging framework of the Daffodil system [16]. While no searcher spent less than 40 minutes on the task, all searchers were asked to stop at the 50 minutes mark. On average they spend about 47 minutes searching, examining documents, and storing results. A combined total of 198 queries were issued by the searchers, and 143 items stored as relevant (an average of 16.5 queries and 11.9 stored items per user).

During their tasks the participants requested search advice a total of 94 times (about 7.8 times per searcher or about once every two queries). After reading the search advice, 62 suggestions were directly followed (i.e. executed semi-automatically) by the participants.

The experiment was accompanied by a two-part questionnaire. The first part of the questionnaire was designed to gather information about the domain and search knowledge of the participants. The second part was given after the participant had concluded the task, and contained four question blocks in which the searchers were asked to judge their experiences with the system.

4.1 Results

Because of the low number of participants the results from the pilot study can only be taken as first indicators. However, some of the results were promising and merit further exploration of the Suggestion Tool. The questions used a seven-point Likert scale from 'fully disagree' (coded as 1) to 'fully agree' (coded as 7). Users were invited to elaborate on their answers, and were asked to give specific examples of problems.

Clarity. Nearly all participants found the suggestions to be easy to understand. The explanations why a specific suggestion might be useful were generally found to be clear (modal score of 7, median score of 6.5). Participants rarely encountered suggestions that they interpreted differently from the intended meaning and where they were surprised by the results of the advice (modal and median score of 6).

A particular problem that was mentioned by several users was the description of suggestion #16, which didn't explain clearly how the list of terms was derived.

Appropriateness. Since the suggestions were selected and ranked according to their similarity to previous users' situations, it was interesting to see if this ranking method arranged the suggestions in a contextually appropriate manner. While the results were only marginally positive (modal score of 4, median score of 5), it turned out that this was mainly the result of the experimental setup. Nearly all problems reported were with very low scoring suggestions, which had been included in the presentation to give the participants a broader range of suggestions to choose from. A threshold to keep out low scoring suggestions is clearly necessary. Suggestions with scores of 0.5 and higher were universally deemed appropriate.

Usefulness. The high scoring suggestions in each situation were found to be generally useful and helpful (median and modal score of 6). Ten of the twelve participants reported that thanks to the suggestions, they had employed tactics and stratagems that they normally wouldn't have used or thought of. All ten stated that they found those helpful and would try to use them independently in future search tasks.

Particular suggestions that were mentioned as helpful were the hints about the different tools from the Daffodil toolbox (thesaurus, classification browser, social network tool), the purpose or even existence of which were not apparent to novice or casual users of the system. The experienced users liked the extraction of common authors, conferences or terms from the result set that was provided by the suggestions. Several mentioned that while they might have eventually used some of the suggested tactics on their own, the advice provided by the ST helped them to realize some of the available opportunities more quickly and to avoid lengthy and ultimately frustrating trial-and-error.

Interface. The visual presentation of the suggestions was generally well-liked, and the participants easily understood the ranking, as well as the icons used by the tool. 75% of the users also preferred that the suggestions were only shown on specific request instead of automatically. The availability indicator proved to be sufficient to notify the user about the existence of new suggestions, while still being unobtrusive enough not to interfere with the original search.

Interaction with the tool proved to be more problematic. Two major flaws were discovered which impaired the initial usability:

1. The icon marking a specific suggestion as "automatically executable" was readily understood. However, because of its similarity to the play button used in many media players, a large number of users tried to execute the suggestion by clicking on the icon (instead of using double-click activation as expected).
2. Although most of the users were willing to judge good suggestion, many had problems finding the option to do this, as it was hidden in the context menu of the suggestion.

To fix these problems, both actions (executing as well as judging) will be made available to the users by easily accessible single-click buttons for each suggestion.

5 Summary and Conclusion

In this paper we have presented a new tool for giving automated advice during the search process. This tool has been successfully integrated into the Digital Library toolbox DAFFODIL, and uses case-base reasoning techniques to find and rank the most appropriate suggestions for the user's current search situation.

A small scale pilot evaluation was conducted, during which 12 users worked with the Suggestion Tool for a total of 560 minutes. The evaluation pointed out a number of possible improvements, but overall user reception of the suggestions was positive. The users found the automated, non-intrusive advice to be helpful, and implemented suggested tactics with success to further their search task. These results confirmed similar findings from [10] and [15].

Several problems that were found during the evaluation were fixed and the list of available suggestions has been extended since and continues to be extended. Further evaluations need to be conducted to improve the ranking of suggestions and determine a cut-off point were the similarity of cases no longer merits their inclusion in the list of possible advices. In addition, an extension of the tool from suggestions of single, isolated tactics and stratagems towards sequences or paths of actions, as originally proposed in [17], is being considered.

References

1. Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1):39–59, 1994.
2. M. J. Bates. Information search tactics. *Journal of the American Society for Information Science*, 30(4):205–214, 1979.
3. M. J. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 13(5):407–424, 1989.
4. Marcia J. Bates. Where should the person stop and the information search interface start? *Information Processing and Management*, 26(5):575–591, 1990.
5. Nicholas J. Belkin. Interaction with texts: Information retrieval as information seeking behavior. In G. Knorz, J. Krause, and C. Womser-Hacker, editors, *Information Retrieval '93. Von der Modellierung zur Anwendung. Proc. d. 1. Tagung Information Retrieval*, pages 55–66, Konstanz, 1993.

6. Nicholas J. Belkin, Colleen Cool, Adelheit Stein, and Ulrich Thiel. Cases, scripts, and information-seeking strategies: On the design of interactive information retrieval systems. *Expert Systems with Applications*, 9:1–30, November 1995.
7. Nicholas J. Belkin, P.G. Marchetti, and Colleen Cool. BRAQUE: Design of an interface to support user interaction in information retrieval. *Information Processing and Management*, 29(3):325–344, 1993.
8. Suresh K. Bhavnani, Karen Drabenstott, and Dragomir Radev. Towards a unified framework of IR tasks and strategies. In *Proceedings of ASIST 2001: 64th Annual Meeting*, pages 340–354, 2001.
9. Giorgio Brajnik, Stefano Mizzaro, and C. Tasso. Evaluating user interfaces to information retrieval systems: A case study on user support. In *Proceedings of the SIGIR 1996*, pages 128–136, 1996.
10. Giorgio Brajnik, Stefano Mizzaro, Carlo Tasso, and Fabio Venuti. Strategic help in user interfaces for information retrieval. *Journal of the American Society for Information Science and Technology*, 53(5):343–358, 2002.
11. Karen M. Drabenstott. Do nondomain experts enlist the strategies of domain experts. *Journal of the American Society for Information Science and Technology*, 54(9):836–854, 2003.
12. Norbert Fuhr, Claus-Peter Klas, André Schaefer, and Peter Mutschke. Daffodil: An integrated desktop for supporting high-level search activities in federated digital libraries. In *Proceedings of the ECDL 2002*, pages 597–612, Heidelberg et al., 2002. Springer.
13. Norbert Gövert, Norbert Fuhr, and Claus-Peter Klas. Daffodil: Distributed agents for user-friendly access of digital libraries. In *Proceedings of the ECDL 2000*, pages 352–355, Heidelberg et al., 2000. Springer.
14. Ingrid Hsieh-Yee. Effects of search experience and subject knowledge on online search behavior: Measuring the search tactics of novice and experienced searchers. *Journal of the American Society for Information Science*, 44(3):161–174, 1993.
15. Bernard J. Jansen. Seeking and implementing automated assistance during the search process. *Information Processing and Management*, 41(4):909–928, 2005.
16. Claus-Peter Klas, Hanne Albrechtsen, Norbert Fuhr, Preben Hansen, Sarantos Kapidakis, László Kovács, Sascha Kriewel, András Micsik, Christos Papatheodorou, Giannis Tsakonas, and Elin Jacob. A logging scheme for comparative digital library evaluation. In *Proceedings of the ECDL 2006*, pages 267–278, Heidelberg et al., 2006. Springer.
17. Sascha Kriewel. Finding and using strategies for search situations in digital libraries. *Bulletin of the IEEE Technical Committee on Digital Libraries*, 2(2), 2006. http://www.ieee-tcdl.org/Bulletin/v2n2/kriewel/kriewel.html.
18. Sascha Kriewel, Claus-Peter Klas, André Schaefer, and Norbert Fuhr. Daffodil - strategic support for user-oriented access to heterogeneous digital libraries. *D-Lib Magazine*, 10(6), June 2004. http://www.dlib.org/dlib/june04/kriewel/06kriewel.html.
19. Barbara M. Wildemuth. The effects of domain knowledge on search tactic formulation. *Journal of the American Society for Information Science and Technology*, 55(3):246–258, 2004.