

A Logging Scheme for Comparative Digital Library Evaluation

Claus-Peter Klas¹, Hanne Albrechtsen², Norbert Fuhr¹, Preben Hansen⁵,
Sarantos Kapidakis⁴, Laszlo Kovacs³, Sascha Kriewel¹, Andras Micsik³,
Christos Papatheodorou⁴, Giannis Tsakonas⁴, and Elin Jacob⁶

¹ University of Duisburg-Essen, Duisburg, Germany

² Institute of Knowledge Sharing, Copenhagen, Denmark

³ MTA SZTAKI, Budapest, Hungary

⁴ Ionian University, Kekryra, Greece

⁵ Swedish Institute of Computer Science, Kista, Sweden

⁶ Indiana University Bloomington, Bloomington, USA

Abstract. Evaluation of digital libraries assesses their effectiveness, quality and overall impact. To facilitate the comparison of different evaluations and to support the re-use of evaluation data, we are proposing a new logging schema. This schema will allow for logging and sharing of a wide array of data about users, systems and their interactions. We discuss the multi-level logging framework presented in [19] and describe how the community can add to and gain from using the framework. The main focus of this paper is the logging of events within digital libraries on a *generalised, conceptual level*, as well as the services based on it. These services will allow diverse digital libraries to store their log data in a common repository using a common format. In addition they provide means for analysis and comparison of search history data.

1 Introduction

Evaluation of digital libraries (DLs) aims at assessing their effectiveness, quality and overall impact. Analysis of transaction logs is one evaluation method that has provided DL stakeholders with substantial input for making managerial decisions and establishing priorities, as well as indicating the need for system enhancements. However, the quantitative nature of this method is often criticised for its inability to provide in-depth information about user interactions with the DL being evaluated. The results of logging studies are often localised and not easily interpretable outside the DL being investigated. The problem of generalisability is compounded by the absence of a standardised logging scheme that could map across the various logging formats being used. The development of such a scheme would facilitate comparisons across DL evaluation activities and provide the means for highlighting critical events in user behaviour and system performance.

The logging scheme proposed in this paper is framed within the DL evaluation activities of the DELOS Network of Excellence. In our whitepaper on DL evaluation [9], we have identified the long-term goal of building a community for DL evaluation. Under the umbrella of an experimental framework that will serve as a theoretical and practical platform for the evaluation of DLs, the proposed logging scheme will allow for meaningful interpretation and comparison of DL transactions. By using this scheme, researchers will be able to extract re-usable data from the results of previous DL evaluations and to identify useful benchmarks, allowing for more efficient and effective design of evaluation studies.

In order to support the re-use of all possible evaluation data, we intend that the proposed scheme will account for all manner of data that can be collected from the user, the system and the user-system interaction. To this end, we are proposing a novel, multi-level logging framework that will provide complete coverage of the different aspects of DL usage. The main focus of this paper is the level of conceptual, generalised user actions, for which we describe the logging scheme in some detail. Based on this specification, we present tools that can help DL stakeholders to analyse the logging data according to their specific interests.

The remainder of this paper is structured as follows: A brief survey of related work on DL logging is given in Section 2. The levels of the proposed logging scheme are presented in Section 3, while Section 4 addresses the events that comprise the conceptual aspects of user actions. Section 5 presents an application of the logging scheme in the DAFFODIL DL. Using researchers as an example, we discuss how the various DL stakeholders can gain from a common logging scheme in Section 6. In conclusion, Section 7 summarises the arguments presented in this paper and outlines future work with logging schemes.

2 Related Work

In DL research, there is a tendency to identify patterns of information searching behaviour through the use of system features such as Boolean operators, use of ranking and sorting mechanisms, identification of the type and nature of queries submitted and analysis of users' distribution of these features [17, 24, 18].

Transaction logs also provide a useful resource for the remote evaluation of web-based information systems due to their ability to record every action that occurs during the user's interaction with a digital resource [6]. They are generally used to gather quantitative data and to optimise the generation of statistical indicators. Logs are widely employed in the DL domain because of the consistency they provide with respect to the conditions of data collection; but logging should be understood as only one aspect of the overall experimental framework, as suggested by [4]. For websites and for bibliographic systems such as OPACs, logs provide dependable and coherent information about the usage, traffic and performance of a given system. However, because logs can only partially reveal both the behaviour of the user and her level of satisfaction [7], richer information is frequently derived by applying other methods of analysis such as sequential pattern analysis and web surveys.

Goncalves and others proposed an XML log standard for digital library logging analysis [12, 11]. They concentrate on high level events that are generated by user actions, and describe events for searching, browsing, updating and storing actions. However, as pointed out by Cooper in [5], for a comprehensive transaction log analysis different sources of log data have to be taken into account. We distinguish between different aspects and levels of logging, as described in Section 3. Each of these levels is supported by a standard XML schema that defines the events of interest according to the special needs of the different stakeholders interested in the logging data.

In addition, we extend the original set of events that have been proposed to more comprehensively support the logging of actions allowed by modern DL services. For the analysis of logged events, we also differentiate between the various stakeholders of a DL system, including, for example, system owners, librarians, endusers, developers and researchers. Each stakeholder group requires a particular view on the logging data in order to address its specific information needs. To support these views on transaction log data and to combine and analyse data from the different aspects of logging, a number of tools is being developed.

There are other efforts to standardise aspects of transaction logs or to provide uniform classifications for their analysis. Yuan and Meadow [27] provide a codification of, among other aspects, the variables of participants in user studies.

3 Levels of Logging

When using transaction logs for evaluation, the main participants under survey are the user and the system, as well as the content that is being searched, read, manipulated, or created. The interaction between the system and the user can be examined and captured at various levels of abstraction.

1. System parameters
2. User-system interaction
 - UI events (keystrokes, mouse movement, etc.)
 - Interface action events (data entry, menu selections, etc.)
 - Service events (use of specific DL services)
 - Conceptual events (generic actions)
3. User behaviour

On a very low level, researchers or developers might be interested in getting information about key presses, mouse clicks and movements, and similar concrete interactive events. These can be grouped together into more meaningful interactive events with specific graphical tools or interface elements, several of which in turn might comprise an interaction with a digital library service like a citation search service. For all three of these levels of abstractions a common logging schema will help in comparison, e.g. of the usability of specific interfaces or services between different digital libraries.

An additional level of abstraction, that will be called the conceptual level, can be seen as an abstraction away from the concrete services of specific digital

library systems towards generalised actions that users can take in using a DL system. Several generic services that could be captured at this level have been proposed in [12] and [11]. Herein we suggest a modified and expanded version of this selection.

By separating the log data into user data, system data and different abstraction levels of interaction – all connected by timestamps and identifiers linking a specific user with each event – the transaction logs can be used for following a users actions chronologically along a specific tier of the model. On the other hand, a comprehensive analysis of what exactly happens during a search action of the user across the various tiers is also possible.

Of course a digital library system might only capture data on a low level of abstraction to minimise the overhead incurred by the logging, and use a mapping of UI events to more abstract events for later analysis.

3.1 System

The system tier of the logging model describes the changes of the system over time, as represented by various parameters of the hardware and software involved in providing the digital library services. These parameters can be captured directly by the backend, usually without involvement of the client software used to access the digital library.

Aspects of the system fall into two groups: static parameters like operating system, available memory, bandwidth or computing power, and dynamic parameters that change over time, usually in reaction to user action, like server load, amount of connected users, network traffic and ping times.

3.2 User Behaviour

At the other end of the spectrum is the tier representing the user and her behaviour, her changing and evolving cognitive model, and her interactions with the environment outside the digital library system. While system behaviour is usually easiest to capture, user behaviour can hardly be captured through transaction logs. Other methods are common within user studies, e.g. video capturing, think aloud studies, search diaries, interviews or questionnaires.

As with systems, aspects describing the user can be grouped into static and dynamic parameters. Static parameters like age, first language, professional or search background, social and organisational environment, usually don't change over the course of one session. Dynamic parameters on the other hand, might include frustration levels, or the user's progression along the stages of her information task.

3.3 User-System Interaction

The logging of the interaction of user and system in distributed systems is complicated by the fact that much of the interaction occurs at the client. Low-level interactive events therefore need to be captured at the site of the user and be transmitted to the server, which introduces difficulties for browser based user

clients. In these cases only higher level interactions with corresponding events on the server side of the digital library might be logged and analysed later [14].

UI Level. On the lowest level of user-system or human computer interaction (HCI), events consist of single input events like keystrokes, mouse clicks or mouse movements. These are of interest for usage and usability studies of systems, and have been analysed and studied in HCI research for a long time [15]. This level corresponds to the keystroke level of the GOMS model [16] and can be combined with low-level events captured by other means (e.g. eye movements).

Interface Level. It is common practice in HCI research [13, 14] to group low-level events into higher-level abstractions for specifying more general models of the user-system interaction. A file selection that incorporates several mouse movements, clicks and textual input is combined into a single, abstract interface action. On this level e.g. the number and kind of interface actions necessary to complete a specific task can be compared between several digital library systems.

Service Level. In a further step of abstraction the service level combines several interface actions into more meaningful services provided by digital library systems. Most DL systems offer a number of different services that support searching and other tasks of the information seeking process, e.g. metadata or fulltext search, search for citations, annotating of documents, services for organising personal collections of DL objects, or for supporting the reviewing of documents. Depending on the tools and options offered by a specific system, it is possible that different combination of interface actions can be combined to use the same service.

Conceptual Level. While the first three levels of user-system interaction combine actions from the level below into a larger, higher-level action, the conceptual level represents an abstraction away from the concrete implementation of specific services, and tries to define generalised types of events. These generalised events or conceptual events represent the various actions that a user might pursue in a digital library system. While this list is probably not comprehensive, care has been taken to describe these actions in a general way that can be applied to most of today's DL systems.

4 Events on the Conceptual Level

On the conceptual level, we have identified several general event types that support comparative evaluation across DLs. These events are partially in line with the events proposed in [12, 11]. They identified some generic events – search, browse, update, store – and some higher concepts – annotate, filtering, recommending, rating, reviewing – which they call transactions. Our focus on the conceptual level represents the centrality of these events for log analysis and interpretation, because they indicate critical aspects of the user's interaction with the DL system and supply valuable data for rich interpretation of user

behaviour. As has been highlighted in other DL logging studies [22], current approaches are often inadequate for capturing complex or abstract actions by the user and are therefore unable to elicit meaningful conclusions.

By logging data about general event types at the concept level, we provide a basis for *comparative evaluation* across DLs.

The event types and event properties that we have identified are neither fixed nor a comprehensive model of user-system interaction and should be viewed as recommendations that can also serve as discussion points. Each event consists of its own set of properties modelled in XML as sub-elements and attributes. Properties that are common for all events are:

- a unique session-id
- start-stop timestamps
- possible errors during the event
- a unique event-id
- a service name
- a cancelation indicator

In addition, each event as described herein also has an event-specific set of properties which are summarized in Table 1. If the collection of additional data about specific events is necessary for a study, it is suggested to extending standard event definitions through reference to an XML namespace that defines the new properties. For example, the standardised search event describes the search condition as a list of terms. However, many DLs allow for more complex query formulations such as Boolean queries, which could be stored in a specific field defined in the extended namespace as a sibling of the list-of-terms input element. Although comparison across DLs will only work on the list-of-terms property defined in the original namespace, researchers who require an extended view of logging events are guaranteed that no information will be lost in the process of comparison.

Search events represent any action of users that involve formulating a query or filter condition that is to be processed by a DL service against a collection. The collection can be the entire document space, already be pre-filtered or even be the result of a previous query. The system response consists of the subset (e.g. in the form of a ranked list) of objects from the initial collection that satisfy the given condition.

Navigate events represent actions that consist of selecting a specific item from a set of items or following a link to its target. This conceptual event includes the use of hyperlinks to navigate within a set of hypertext documents, but also navigating through a representation of a social network (e.g. of co-authors).

Inspect events capture the user actions of accessing a detailed view of a single object, like the metadata or fulltext of a result document. Similarly, looking up a definition of a term in an encyclopedia or dictionary, or semantic information from a thesaurus is seen as an *inspection* of this term.

Display events describe specific visualisations of DL objects. While the actual content of the presented information does not necessarily change, the change in view on this information is represented by a display event. This conceptual event encompasses actions ranging from a simple resorting of a list

Table 1. Subelements and properties of events

Search	query or filter condition, collection to be searched, system response
Navigate	link to be followed, current collection, system response
Inspect	object to be inspected, system response
Display	collection to be (re-) displayed, visual transformation or visualisation to be applied, sort criteria (optional)
Browse	collection to be browsed, method and dimensions of browsing, direction and distance that the view point is moved
Store	set of DL objects to be stored, target location, method of storage
Annotate	document (optionally part of) to be annotated, type and content of the annotation (may be one or more other DL objects)
Author	the new document, optionally identifier of changed document
Help	help request (optional), type and content of system suggestion
Communicate	type, content and recipients of message

of objects or the presentation of a set of terms in form of a tag cloud, up to complex visual representations of abstract information. In the interest of comparability, the use of a standardised taxonomy or classification of visualisation techniques is proposed, e.g. based on Shneiderman's task by data type taxonomy [25] or the classification of visualisation techniques in [3].

Browse events describe user actions that involve changing the view point on a set of DL objects without changing the visualisation or navigating to a different set of items. Typically these actions will involve scrolling in one or more dimension, using sliders to zoom in or out of a visualisation, or “thumbing” through a document. If the original set of documents has been split into several chunks, browsing might also describe moving from one chunk of the document set to the next or previous one.

Store events are actions of the user or the system that create a permanent or temporary copy of a digital library objects. This might be a digital copy to a clipboard for temporary storage during a search session or to a more permanent location either within the DL system or outside (e.g. on an optical medium or in a web storage). Storing a digital library object can also mean converting from digital to physical form (printing a document), or exporting to a special citation format.

Annotate events cover any user action that adds additional information to an existing DL object, which may be user-specific, shared among a group of collaborators, or visible system wide. The general annotation event includes marking entire documents or specific parts, adding ratings, tags [10] or textual comments like reviews or summaries to an object, or linking two or more DL objects.

Author events describe the creation of a new DL object or direct editing (not annotating) of an existing one. Authoring a document can include writing

a review or another type of textual annotation, or can be part of creating a completely new document if supported by the digital library system.

Help events from the user's point of view can be of a passive nature, where the system provides unprompted suggestions to the user, or of a more interactive nature. In the latter case the user explicitly requests help, suggestions or recommendations about a specific or general topic, and the system generates a response to that. Unprompted help can take the form of recommendations about content, users or specific actions, or provide explanation about functions or system activity.

Communicate events capture events that occur during the communication between two or more users of the digital library system. The communication can be textual or include other media. This general event includes the direct sharing of a DL object with another user and sending messages to other users by using instant messaging, message boards or e-mail components of the system. More technically sophisticated systems might allow for sending voice or video messages between users as well. Digital library services for collaboration will typically also include means of communication for managing the collaboration.

Some of the events contain the actual digital library objects either in form of object identification like DOI, URL or URI. If such a unique identification is not possible, the main or all metadata fields should be provided to distinguish the objects.

5 Logging in Daffodil

DAFFODIL¹ [23, 21, 8] is a virtual DL targeted at strategic support of users during the information search process. For searching, exploring and managing DL objects DAFFODIL provides information seeking patterns that can be customised by the user for searching over a federation of heterogeneous digital libraries. Searching with DAFFODIL makes a broad range of information sources easily accessible and enables quick access to a rich information space.

5.1 Logging

DAFFODIL is an application consisting of a graphical user interface and back-end services written in Java. This makes logging either UI events within the client or back-end event triggered by user actions or by the system a much simpler process than for web-based DL systems. The logging service itself, depicted in Figure 1, is simplistic: the DAFFODIL user interface client and each DAFFODIL back-end service can send an event to the event logging service, which then stores the event.

Currently, DAFFODIL handles over 40 different events. The main groups of events are search, navigate and browse events; result events are generated by each

¹ <http://www.daffodil.de>

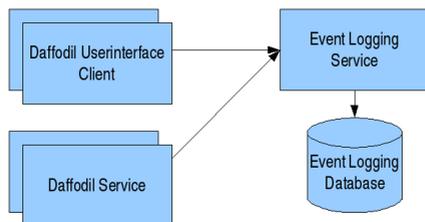


Fig. 1. Logging Service Model

of the system services (e.g., the thesaurus, the journal and conference browser, or the main search tool). The personal library supports store events as well as events involving annotation or authoring of objects.

Through the DAFFODIL service for log schema conversion, we are able to provide more than 100 MB of log data in the format of the proposed XML schema. The data will be anonymised and soon be made accessible for comparative analysis.

6 Analysis of Log Events

In order to analyse the logged events, we have assumed that different stakeholders need different views of the logging data; thus, a variety of analysis tools is required. We have identified the following DL stakeholders: *system owners*, *content providers*, *system administration*, *librarians*, *developers*, *scientific researchers*, and *end-users* of a DL.

A number of tools for facilitating analysis on log data in the new scheme have already been implemented. The example statistics shown in Figures 2 and 3 were produced with the help of these tools.

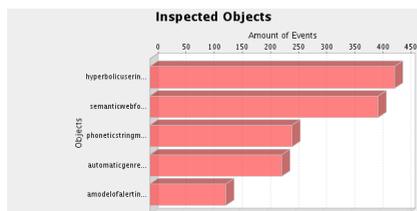


Fig. 2. Inspected Objects

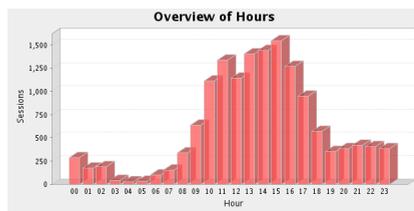


Fig. 3. Usage per hour

Having an experimental framework is a special boon for **researchers** in the field of digital libraries. With DAFFODIL and its use of the standardised logging scheme, a baseline system and a powerful toolbox for evaluation work can be provided. For research on higher concepts, the researchers do not have to develop

or implement a complete setting, but can re-use the framework and build on it. At the system level, the efficiency of algorithms or the appropriateness of DL architectures can be evaluated and compared. HCI research on the usability of digital library interfaces as well can benefit from a framework that provides standardised, comparable logging data.

The major research focus of the DAFFODIL project has always been to place the user in the center of digital library and information retrieval research. DAFFODIL is based on the information search model by Marcia Bates [2, 1] which classifies search activities of users as moves, tactics, stratagems and strategies. The proposed logging scheme and the concept level is a natural extension of this original aim, as it provides a method to analyse the sequences of moves and tactics. As Wildemuth stated in [26]:

While significant work has examined the individual moves that searchers make (Bates, 1979, 1987; Fidel, 1985), it is equally important to examine the sequences of moves made by searchers in order to understand the cognitive processes they use in formulating and reformulating their searches.

With regard to this goal, the logging scheme allows for studying the usage of the diverse services; furthermore, the whole search process/sequence can be analysed, from the initial formulation of a search query to its conclusion with the storage of DL objects for further use.

In addition to capturing and understanding users' search activities through analysis of logging data, recommendations can be made for re-use of search patterns. Kriewel [20] suggests that search paths discovered in analysis of logging data can be used as a basis for suggesting potential search steps to other users.

Of course the vertical look at the levels can also be under examination.

7 Summary and Outlook

In this paper, we have presented the first efforts to develop a standardised experimental framework for digital library evaluation. If the various DL stakeholders can form a community and agree upon models, dimensions, definitions and common sense criteria for the evaluation of DL systems, the process of evaluation will gain impetus among DL researchers. As an experimental platform for evaluation both within and across DL systems, application of the DAFFODIL framework can substantially advance this research area, since it currently provides functions and services that are based on a solid theoretical foundation and well known models. The proposed logging scheme is a first step intended to encourage evaluation of individual systems as well as comparisons across systems.

Most evaluation techniques require a great deal of preparation and effort and are thus not easily replicated. In case of online DL systems, this means that the results of an evaluation often reflect a past snapshot of the system. It is necessary to find ways for continuous, cost effective and (more or less) automated evaluation of digital libraries. The suggested logging scheme is a first

step in this direction. Our group aims at establishing a community forum for evaluators in order to promote the propagation of various tools and approaches, and the exchange of experience.

As part of the effort to encourage a community forum for researchers interested in DL evaluation, we have published documentation for the logging scheme on the forum website². The log analysis tools will follow soon, as they are under preparation. As a further step, large amounts of anonymised logging data from two DL services will be made available. Such primary data will help the community to improve the application of transaction logging and to compare and experiment with sample data. In the long term, we envision that this effort will evolve into a primary data repository, providing help for evaluators who want to find similar scenarios together with logging data. In other fields of research, such primary data repositories are already well established and play an important role in the conduct of research. (In fact, the provision of primary data frequently counts as a publication, creating further incentives for this kind of work.) An infrastructure of independent evaluator services, primary data repositories, logging tools and on-line questionnaires may provide computer-based support for some of the cost-effective and time-consuming tasks in evaluation and the community will gain sustainability. So we ask the community to form, discuss and agree upon a schema, to add data and service in order make a step forward in DL evaluation.

References

- [1] M. J. Bates. Idea tactics. *JASIS*, 30(5):280–289, 1979.
- [2] M. J. Bates. Information search tactics. *JASIS*, 30(4):205–214, 1979.
- [3] E. Bertini, T. Catarci, S. Kimani, and L. D. Bello. Visualization in digital libraries. In *From Integrated Publication and Information Systems to Virtual Information and Knowledge Environments*, pages 183–196, 2005.
- [4] J. Bertot, C. McClure, W. Moen, and J. Rubin. Web usage statistics: measurement issues and analytical techniques. *Government Information Quarterly*, 14(4):373–395, 1997.
- [5] M. D. Cooper. Design considerations in instrumenting and monitoring web-based information retrieval systems. *JASIS*, 49(10):903–919, 1998.
- [6] M. D. Cooper. Usage patterns of a web-based library catalog. *JASIST*, 52(2):137–148, 2001.
- [7] D. T. Covey. Usage and usability assessment: Library practices and concerns. Technical report, Digital Library Federation, 2002.
- [8] N. Fuhr, C.-P. Klas, A. Schaefer, and P. Mutschke. Daffodil: An integrated desktop for supporting high-level search activities in federated digital libraries. In *Research and Advanced Technology for Digital Libraries. 6th European Conference, ECDL 2002*, pages 597–612. Springer, 2002.
- [9] N. Fuhr, G. Tsakonas, T. Aalberg, M. Agosti, P. Hansen, S. Kapidakis, C.-P. Klas, L. Kovcs, M. Landoni, A. Micsik, C. Papatheodorou, C. Peters, and I. Solvberg. Evaluation of digital libraries, 2006.

² http://www.is.informatik.uni-duisburg.de/wiki/index.php/JPA_2_-_WP7

- [10] S. A. Golder and B. A. Huberman. The structure of collaborative tagging systems. Technical report, Information Dynamics Labs, HP Labs, 2005.
- [11] M. A. Goncalves, E. A. Fox, L. Cassel, A. Krowne, U. Ravindranathan, G. Panchanathan, and F. Jagodzinski. Standards, mark-up, and metadata: The xml log standard for digital libraries: analysis, evolution, and deployment. In *Proceedings of the third ACM/IEEE-CS joint conference on Digital libraries*, 2003.
- [12] M. A. Goncalves, R. Shen, E. A. Fox, M. F. Ali, and M. Luo. An xml log standard and tool for digital library logging analysis. In *Agosti, Maristella (ed.) et al., Proc. of 6th ECDL*, pages 129–143, 2002.
- [13] M. A. Hearst and M. Y. Ivory. The state of the art in automating usability evaluation of user interfaces. pages 470–516, 2001.
- [14] J. Helms, D. Neale, and P. I. and J.M. Carroll. Data logging: Higher-level capturing and multi-level abstracting of user activities. In *Proceedings of the 40th annual meeting of the Human Factors and Ergonomics Society*, 2000.
- [15] D. M. Hilbert and D. F. Redmiles. Extracting usability information from user interface events. *ACM Computing Surveys*, 32(4):384–421, 2000.
- [16] B. E. John and D. E. Kieras. Using goms for user interface design and evaluation: Which technique? acm trans. *ACM Trans. Comput.-Hum. Interact.* 3, 1996.
- [17] S. Jones, S. J. Cunningham, R. McNab, and S. Boddie. A transaction log analysis of a digital library. *International Journal on Digital Libraries*, 3(2):152–169, 2000.
- [18] H. Ke, R. Kwakkelaar, T. Tai, and L. Chen. Exploring behavior of e-journal users in science and technology: Transaction log analysis of elsevier’s sciencedirect onsite in taiwan. *Library & Information Science Research*, 24(3):265–291, 2002.
- [19] C.-P. Klas, H. Albrechtsen, N. Fuhr, P. Hansen, E. Jacob, S. Kapidakis, L. Kovacs, S. Kriewel, A. Micsik, C. Papatheodorou, and G. Tsakonas. An Experimental Framework for Comparative Digital Library Evaluation: The Logging Scheme. In *JCDL, Short Paper*, 2006.
- [20] S. Kriewel. Finding and using strategies for search situations in digital libraries. Bulletin of the IEEE Technical Committee on Digital Libraries (to appear), 2005.
- [21] S. Kriewel, C.-P. Klas, A. Schaefer, and N. Fuhr. Daffodil - strategic support for user-oriented access to heterogeneous digital libraries. *D-Lib Magazine*, 10(6), June 2004.
- [22] B. Pan. Capturing users behavior in the national science digital library (nsdl). Technical report, NSDL, 2003.
- [23] A. Schaefer, M. Jordan, C.-P. Klas, and N. Fuhr. Active support for query formulation in virtual digital libraries: A case study with DAFFODIL. In *Proc. of 7th ECDL*, 2005.
- [24] M. Sfakakis and S. Kapidakis. *User Behavior Tendencies on Data Collections in a Digital Library*, volume 2458 of *Lecture Notes In Computer Science*, pages 550–559. Springer-Verlag, Berlin; Heidelberg, 2002.
- [25] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. Technical Report CS-TR-3665, University of Maryland, Department of Computer Science, July 1996.
- [26] B. M. Wildemuth. The effects of domain knowledge on search tactic formulation. *JASIST*, 55(3):246–258, 2004.
- [27] W. Yuan and C. T. Meadow. A study of the use of variables in information retrieval user studies. *JASIS*, 50(2):140–150, 1999.