1a. Project Title

Detecting relations between search queries and search results to improve access to image repositories.

1b. Project Acronym

-

1c. Principal Investigator

Jacco van Ossenbruggen

1d. Renewed Application

Based on the VENI proposal submitted by Vera Hollink on 7 January 2010 (under review).

2a. Scientifc Summary

Finding images remains a difficult task where users modify their queries several times in order to fulfill their information needs. Our goal is to make search more efficient by developing methods to assist users during query modification.

Existing techniques do not suffice as a base for developing search support. Log analysis uses users' queries and the results selected, but not the *meaning* of the queries. As a result, log analysis yields statistics on the overlap in terms between queries, but not their semantic relations. User studies do reveal semantic relations, but are based on limited numbers of users, making statistics about the use of the various relations unreliable.

We propose to combine statistical search log information with semantic information. The first research challenge is to elevate queries from meaningless strings to entities with well defined properties and explicit relations to other entities, by relating the queries to entities from publicly available linked data sources. Statistical methods will be investigated to infer *semantic modification patterns* from frequently occurring relations between queries. The second challenge is to employ these patterns to automatically improve search support.

Offline experiments using the search logs of four large image providers will measure the predictive value of the modification patterns. The impact of the developed search facilities will be evaluated by incorporating them in a live search engine and comparing the search of users who have access to the extra search facilities to users using the original version of the search engine.

2b Abstract for laypersons (in Dutch)

Het zoeken naar afbeeldingen en andere niet-tekstuele documenten is voor veel mensen nog steeds moeilijk. Onderzoek wijst zelfs uit dat gebruikers die zoeken naar plaatjes gemiddeld 80% meer zoekpogingen nodig hebben dan gebuikers die zoeken naar tekst. Dit terwijl het maatschappelijk belang van het efficient kunnen vinden van niet-tekstuele documenten snel toeneemt. Professionele gebruikers in de nieuws en media sector zoeken dagelijks naar afbeeldingen in grote online collecties. In de wetenschap neemt de hoeveelheid data in niet tekstuele vorm snel toe (dankzij technieken als fMRI in de geneeskunde, telescoop data in astronomie, 3D modellen in de moleculaire biologie, etc). Amateur fotografen konden vroeger slechts met grote moeite die ene familiefoto vinden in een stapel schoenendozen. Nu hebben veel gebruikers veel meer foto's digitaal beschikbaar, op eigen harddisks of op websites zoals flickr.com. Het vinden van die ene foto is er echter niet makkelijker op geworden: vaak zelfs moeilijker! Er is dus een grote noodzaak om 1) te begrijpen hoe mensen zoeken naar afbeeldingen, en waarom ze wel of niet vinden wat ze zoeken, en om 2) technieken te ontwikkelen die het zoeken naar afbeeldingen beter ondersteunen.

Er zijn momenteel twee gangbare manieren om te onderzoeken hoe mensen zoeken, en beide aanpakken hebben inherente nadelen die een werkelijke doorbraak voorkomen.

Een mogelijkheid is om gedrag van individuele gebruikers te onderzoeken, met als voordeel dat dit inzicht kan geven in wat mensen precies *bedoelen* met een zoekterm, en *waarom* gebruikers het resultaat wel of niet relevant vinden. Een groot nadeel is dat de definitie van relevantie verschilt van gebruiker tot gebruiker, en van domein tot domein. Hierdoor is het onmogelijk is om voldoende gebruikers en domeinen te onderzoeken om tot statistisch betrouwbare voorspellingen te komen.

Een andere mogelijkheid die nu gangbaar is, is het onderzoeken van "logfiles", grote bestanden waarin zoekmachines bijhouden waar gebruikers op zoeken en wat ze vinden. Een voordeel van deze aanpak is dat er genoeg data beschikbaar is om statistisch betrouwbare uitspraken te doen. Een groot nadeel is dat de huidige technieken om dit soort logs te analyseren zich voornamelijk op de *syntaxis* van de zoektermen richten, en geen rekening houden met de *betekenis* die gebruikers hechten aan zoektermen.

Ons voorstel is om de nadelen van beide methoden te omzeilen door zoektermen uit logfiles te koppelen aan betekenisvolle definities in thesauri en andere gestructureerde data bronnen (die in steeds grote aantallen vrij op het web verkrijgbaar zijn). Doordat deze bronnen ook de onderlinge relaties tussen concepten expliciet coderen, wordt het mogelijk inhoudelijke relaties te vinden tussen twee of meer verschillende zoektermen. Veel termen zijn op vele manieren gerelateerd, terwijl maar enkele relaties voor een gebruiker relevant zullen zijn. Met dit onderzoek willen we, door relaties uit zeer veel websessies te analyseren, kunnen voorspellen welke relaties relevant zijn om te kunnen vinden wat een gebruiker zoekt. Op basis van dit soort voorspellingen kunnen zoekmachines verbeterd worden, bijvoorbeeld door de gebruiker betere zoektermen te suggereren, of door *inhoudelijk* gerelateerde zoekresultaten toe te voegen aan de nu gebruikelijke *syntactisch* gerelateerde resultaten.

2c Keywords

web usage mining, query modification, search support

3 Classication

Most relevant NOAG-ict theme: Data explosion

4 Composition of the Research Team

Name	Area of expertise	Affiliation
Dr. Jacco van Ossenbruggen	Semantic Web	CWI, VU
Dr. Vera Hollink (position to be funded)	Information Retrieval, User Modeling	CWI
Prof.Dr. Arjen P. de Vries	Information Retrieval	CWI, TU Delft
Prof.Dr. Lynda Hardman	Multimedia Interaction	CWI, UvA

The Interactive Information Access Group of the Centrum Wiskunde en Informatica consists of leading researchers including Dr. Van Ossenbruggen, Prof.Dr.Ir. De Vries, and Prof.Dr. Hardman. The group combines experience on (image) information retrieval with research on using semantics for information disclosure. In the context of various international projects and networks, such as Vitalas, EuropeanaConnect, and PetaMedia, the research group has established several lasting collaborations with organizations that maintain image search engines. These collaborations enable us to work with the search log data of large numbers of real users. Moreover, contacts within these institutes provide a platform to discuss the demands of real image providers and for getting feedback on our solutions. On the topic of using semantics in search the group closely collaborates with the VU Web and Media Group. This collaboration has yielded an award-winning prototype [26], which will serve as a platform on which our results will be evaluated.

5 Research School

SIKS

6a Description of the Proposed Research

Image search engines allow users to search in large image collections. Despite considerable improvements in recent years in the area of content- and annotation-based image retrieval, finding images remains a challenging task. Users searching for images need to modify their queries several times to fulfill their information needs: on average, users searching for images need 20% more search iterations than users searching for audio and video content [28] and even 80% more iterations than users searching for textual content [18]. Apparently, finding effective queries is a substantial component of image search. The aim of the proposed research is to develop methods to assist users with query modification, making image search as a whole more efficient. We focus on *information gathering* tasks [1, 6, 21, 23, 27]: collecting multiple pieces of information around a single topic. Compared to other types of search tasks, information gathering tasks occur most frequently, take up most of the users' time and involve most query modifications [1, 21].

To support query modification we first need to understand how users modify their queries and how these modifications help to fulfill their information needs. This gives rise to the following research questions:

RQ 1 How to identify meaningful patterns of relations between sequences of queries that are entered in a search session and between queries and selected images?

RQ 2 How to apply query modification patterns to improve the support that image search engines offer to their users?

Limitations of current approaches

Existing approaches for analyzing query modification behavior cannot identify semantic modification patterns. Automatic analyses of search logs are purely statistical, using only queries that users have entered and clicks they made on search results [16]. Studies employing this type of analysis have classified query modifications based on the overlap between terms [4, 5, 9, 14, 17, 19, 22, 25, 29]. Term-based methods can only classify pairs of queries that have at least one term in common and, therefore, cannot determine relations between queries that are semantically related but share no terms, such as `Boris Becker` and `Andre Agassi`. Furthermore, these methods recognize the modification of `Venus Williams` to `Serena Williams` as a case in which a term has been substituted, but not as relation between two tennis-players or sisters.

In the field of human-computer interaction, a popular means for researching search behavior is to carry out in-depth studies of the search behavior of a small number of users (e.g. [8, 15]). The (modification) behavior of the users is examined in laboratory experiments or field studies and the users' motivations for the various search steps are revealed through diaries or interviews [20]. In contrast to log analysis, this type of research can reveal semantic relations between queries. However, it is necessarily restricted to a small number of users, which makes statistics about the use of the various types of relations unreliable. Moreover, user studies are costly and time-consuming.

Overview of the proposed research

We propose to combine the strengths of the two existing approaches, automatically identifying semantic relations and employing log statistics to determine the importance of the various types of relations. Search logs are automatically collected when users use a search engine. Large sources of semantic information have recently become freely available in the form of *linked open data* [2, 3]. Although the potential of these two data sources is widely acknowledged, their combination is largely unexplored.

We will match the queries that users enter consecutively during a search session to entities in linked data, thus elevating the queries from meaningless strings to entities of which the semantic relations to other entities are known. Links between the entities will be exploited to find semantic relations between the queries. Comparing statistics about the relations between large numbers of query pairs enables us to identify semantic modification patterns. In the same way, we will use linked data to establish semantic relations between queries that users enter in the beginning of a session and the images that they select later on.

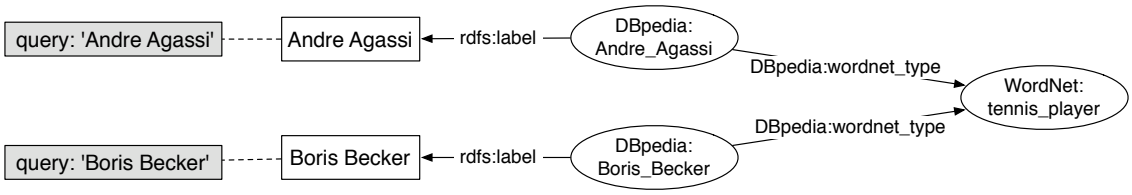

Figure 1: Example application of our method for finding semantic relations between queries: a relation between queries `Andre Agassi` and `Boris Becker` is that they both match DBpedia entities that are of WordNet type `tennis_player`.

Data sets used

We will use the search log data of four image search engines: a Dutch and a Belgian vendor of historic and news photos, a web site of a Dutch art museum and a portal providing access to many cultural collections in Europe. These institutes have made their log data available to the research group as part of longer term collaborations. Having search logs from various partners provides us with the unique opportunity to compare modification patterns in different data sets and determine which patterns are valid across different types of image search and which differ among search engines and user populations.

The linked data that we will use is part of the publicly available data cloud collected in the Linking Open Data Project [24]. According to the latest estimates this data cloud connects over a hundred data sources, in total consisting of over 4.7 billion RDF triples, where each triple defines a property of an entity or a relation between two entities [3]. As the data cloud continues to grow, the methods that we develop will be able to handle more and more different queries.

Approach to answer RQ1: identifying meaningful patterns

Using both log data and linked data, the proposed method matches queries to entities in linked data by examining the similarity between the queries and the labels of the entities in the linked data. To find the relation between the queries, it searches the shortest series of links that connect the entities. Graph search techniques will be employed to reduce the complexity of this shortest path problem. For example we match the query `Andre Agassi` to the concept `http://dbpedia.org/resource/Andre_Agassi` and the query `Boris Becker` to `http://dbpedia.org/resource/Boris_Becker`. Following the relations in the linked data, we can determine relations between the two concepts, such as that both are men and both are tennis players. This process is illustrated in Figure 1. Often multiple relations exist between two entities. Relations that are likely to be important for users are those that occur significantly more frequently between consecutive query pairs than between query pairs from different search sessions, as we will determine through statistical analysis. Semantic modification patterns will be identified by extracting important relations that occur in many search sessions. In preliminary experiments this method revealed many interesting patterns [12].

Approach to answer RQ 2: improving support

Various authors have advocated the use of term-based modification patterns for enhancing search support, in particular through suggestions for follow-up queries [14, 17, 25], but so far this has not been realized. We will take research into query modification a step further by investigating how modification patterns can be applied to improve support during various phases of information gathering tasks and evaluating the support in user studies:

- In the first phase of the search the user (re)formulates a query, translating the information needs into a textual string. We aim to reduce the number of search iterations by offering suggestions for follow-up queries. Preliminary results have demonstrated that semantic modification patterns can be successfully applied for query suggestion [13]. In contrast to existing suggestion methods based on log data, the pattern-based approach can also generate suggestions for first-time queries. Moreover,

the patterns provide explanations of the relations between the query and the suggestions, which in other domains has shown to increase the users' acceptance of system-generated recommendations [7].

- During the second phase the search engine retrieves a set of images that match the user's query. We will investigate under what conditions relations between queries and selected results are usable for predicting from a user's initial queries the images that s/he will select later in the session. The predicted images will be added to the results of the initial queries with the aim of increasing recall and improving the ranking of the search results.

- The final phase involves presenting the search results to the user. We will use modification patterns to explain how the search results retrieved relate to the user's query and to group results by the relations they bear to the query. We expect that this will simplify the selection of relevant search results for the users, shortening the time they need to assess whether results are relevant and reducing the number of clicks on irrelevant results.

We will personalize the search support for all three search phases by adapting it to the characteristics of specific user groups. Clustering techniques, such as those in [30], will be used to find groups of users with similar modifications patterns. This enables us to base the search facilities on patterns of relevant user clusters instead of patterns of the whole user population, which can reduce the users' search efforts even further.

6b. Application Impact

As increasing numbers of images are becoming available in digital image repositories, methods for efficient disclosure of the repositories are urgently needed. The proposed research contributes to the disclosure of digital image collections by making use of the feedback that is implicitly provided when people are using the system. In order to carry out our studies we need data sets and log files, which are available for cultural heritage archives and commercial image providers. However, the proposed methods are by no means restricted to these domains: semantic modification patterns have the potential to enhance search in a wide variety of visual collections, in particular where domain knowledge is available in the form of thesauri.

Cultural heritage archives

At present almost all cultural heritage institutes (museums, archives, libraries) have digitized or are in the process of digitizing the contents of their paper archives. The images are often annotated with terms from domain-specific thesauri. In various large scale projects collections and thesauri of different archives are being connected to each other and to general purpose ontologies, allowing users to search through even larger image collections. The collections are queried by both amateurs and professionals users, such as art historians and museum curators. Queries of professional users tend to be broad and complex, demanding more advanced search techniques than the simple keyword matching offered by most existing search facilities [1, 11].

Our research provides the technologies needed to develop search facilities that help users to answer complex questions faster and more easily. Suggestions for related concepts help users to broaden their search. Explicit explanations for the relations between queries and search results allow users to recognize the relevance of non-obvious results. The specialized thesauri available in the domain, linked with general purpose lexical resources such as WordNet, enable us to provide support for both queries with domain-specific terms and queries with layman terminology.

We apply our methods to the search logs of the search engines of two cultural heritage archives that have made their log data available to the research group: a Dutch art museum and a portal providing access to many European cultural collections. If the results of the evaluation of the created search support is positive, we expect, based on our collaborations with the institutes, that the institutes will incorporate the support in their search engines.

Commercial image providers

Commercial image providers offer vast collections of stock images related to a multitude of topics, such as historical events, news and sports. New content is constantly added to the collections and new topics are emerging all the time. Users have diverse backgrounds and purchase images for a variety of applications. As a result, the system is constantly confronted with users with new information needs, formulating queries that not have been entered before and searching for newly added images.

The dynamic nature of the collection makes the domain unsuitable for search facilities that rely on statistics about the use of individual images and queries. The semantic modification patterns that we identify describe the search behavior of the users on a higher level of abstraction, which makes search facilities based on these patterns robust against changes in content and queries. Personalization techniques that we develop enable us to deal with the diversity in the user population.

Search logs of a vendor of historic photos and a vendor of news photos will be used to test our techniques in the domain of commercial image providers. To these companies the developed techniques are of direct commercial value, as they enable their customers to search the collections more effectively.

Other potential application areas

Medical image collections contain a wealth of information relevant for health-care professionals as well as patients. The collections are growing rapidly due to images collected through daily clinical routines, which heightens the need for effective search support. The medical domain is characterized by a large number of synonyms and a large diversity between terms used by laypeople and medical specialists. As image collections are usually annotated with terms from only one vocabulary, queries using other terminologies often yield no results. Another difficulty is that a dense network of relations exists between medical concepts, many of which are not immediately clear to non-experts. Not showing images annotated with related concepts deprives the user of many potentially useful results. On the other hand, showing all related images makes it difficult for users to understand why the images presented are relevant. Medical ontologies combined with automatically collected log data from search engines for medical images enable the implementation of search facilities based on semantic modification patterns. These facilities provide adequate solutions for the search problems in this domain. Through medical ontologies we can match laypeople's terms with highly specialized terms, enabling us to generate query suggestions that help users to find terms that match the vocabulary of the collection. Furthermore, our methods provide explanations for relations between the user's query and the images that are retrieved.

The rise of digital cameras and affordable storage solutions has dramatically increased the amount of image material in **personal picture collections**. The images in these collections are usually sparsely annotated and the annotations that exist use inconsistent terminologies. One collection often contains many images from the same event or on the same topic. Users searching in their personal collections often want all images on one topic, demanding high recall. This is difficult with the existing search engines as they are mostly focused on precision. The methods that we develop in this project can enhance recall as they facilitate finding images annotated with a terminology different from the user's query and broadening of search to closely related topics.

The techniques that we develop can also be generalized to other media, such as **video**. Also in video collections, searchers experience considerable difficulties formulating adequate queries and need to try various queries before they find what they need [18]. We provide methods to analyze these query modifications and use this information to assist users during their search.

7 Project Planning

During the first six months of the project we will develop the core method for finding semantic relations between consecutive search queries. In the second period of six months we will investigate methods to apply semantic modification patterns for query suggestion. We will evaluate the various approaches using the log data of the four image search engines. At this stage the suggestions will be evaluated offline by measuring how often the next query of a user in the log file is among the suggested queries [10].

The second year will be devoted to identifying and exploiting relations between queries and selected images. We will use terms from the image annotations to match images to concepts in linked data sources. Patterns of these relations will be used to predict which images will be relevant for later queries and to add these to the result lists. The effects of the added images on precision and recall will again be assessed offline.

In the third year we will focus on employing the patterns to improve result presentation. Moreover, all search facilities developed so far will be subjected to a user study. For this, we will make use of an existing search engine for cultural heritage content, which is used by a considerable number of real users. We will incorporate the various facilities into the search engine and offer them as extra support to a random selection of the search engine's users. Differences between the search behavior of users with and without access to the various facilities will allow us to accurately measure the added value of each type of search support.

In the end we expect that our methods will improve the efficiency and efficacy of image search by significantly decreasing the number of iterations users need to fulfill their information needs and increasing the number of relevant images found.

Time line

Months	Topic	Deliverables
1-6	method for identifying semantic modification patterns	conference paper
7-12	employing patterns for query suggestion	journal paper
13-18	method for identifying relations between queries and search results	conference paper
19-24	employing patterns to improve recall	journal paper
25-30	employing patterns to improve search result presentation	conference paper
31-36	large scale user evaluation	conference and journal paper

8 Expected Use of Instrumentation

The project will be executed on CWI's available infrastructure.

9a Literature

Five publications relevant to the research proposal:

- Hollink, Tsikrika and De Vries. The semantics of query modification. *9th International Conference on Adaptivity, Personalization and Fusion of Heterogeneous Information (RIAO 2010)* , Paris, 28-30 April 2010. Accepted for publication.
- Hollink, Van Someren and Wielinga. A Semi-Automatic Usage-Based Method for Improving Hyperlink Descriptions in Menus. *International Journal of Human-Computer Studies* 67, pages 366-381 , 2009.
- Schreiber, Amin, Aroyo, Van Assem, De Boer, Hardman, Hildebrand, Van Ossenbruggen, Omelayenko and Tordai: Semantic Annotation and Search of Cultural-Heritage Collections: The MultimediaN E-Culture Demonstrator. *Journal of Web Semantics* (volume 6, number 4, pages 243 - 249), November 2008
- Tsikrika, Diou, De Vries and Delopoulos. Image annotation using clickthrough data. *Proceedings of International Conference on Video and Image Retrieval* (CIVR), 2009, July 8 - 10, Santorini, Greece. ACM, 2009
- Wang, De Vries and Reinders. Unified relevance models for rating prediction in collaborative filtering *ACM Transactions on Information Systems*, Vol. 26, Nr. 3, June 2008, pp. 1 - 42.

9b References

[1] Alia Amin, Jacco van Ossenbruggen, Lynda Hardman, and Annelies van Nispen. Understanding cultural heritage experts' information seeking needs. In *Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital libraries, Pittsburgh PA, USA*, pages 39–47, 2008.

[2] Tim Berners-Lee. Linked data: Design issues. http://www.w3.org/DesignIssues/LinkedData.html, 2006. last accessed February 13, 2010.

[3] Christian Bizer, Tom Heath, and Tim Berners-Lee. Linked data - the story so far. *International Journal on Semantic Web and Information Systems, Special Issue on Linked Data*, 5(3):1–22, 2009.

[4] Paolo Boldi, Francesco Bonchi, Carlos Castillo, and Sebastiano Vigna. From 'dango' to 'japanese cakes': query reformulation models and patterns. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, Milan, Italy*, pages 183–190, 2009.

[5] Peter Bruza and Simon Dennis. Query reformulation on the internet: empirical data and the hyperindex search engine. In *Proceedings of the RIAO'97 Conference on Computer-Assisted Searching on the Internet, Montreal, Canada*, pages 488–499, 1997.

[6] Chun Wei Choo, Brian Detlor, and Don Turnbull. Working the web: an empirical model of web use. In *Proceedings of the 33rd Hawaii International Conference on System Sciences-Volume 7*, 2000.

[7] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. The effects of transparency on trust in and acceptance of a content-based art recommender. *User Modeling and User-Adapted Interaction*, 18(5):455–496, 2008.

[8] Efthimis N. Efthimiadis. Interactive query expansion: a user-based evaluation in a relevance feedback environment. *Journal of the American Society for Information Science*, 51(11):989–1003, 2000.

[9] Daqing He, Ayse Göker, and David J. Harper. Combining evidence for automatic web session identification. *Information Processing and Management*, 38(5):727–742, 2002.

[10] Qi He, Daxin Jiang, Zhen Liao, Steven C. H. Hoi, Kuiyu Chang, Ee-Peng Lim, and Hang Li. Web query recommendation via sequential query prediction. In *Proceedings of the 25th IEEE International Conference on Data Engineering, Shanghai, China*, pages 1443–1454, 2009.

[11] Michiel Hildebrand, Jacco Van Ossenbruggen, Lynda Hardman, Jan Wielemaker, and Guus Schreiber. Searching in semantically rich linked data: a case study in cultural heritage. *Submitted for publication*.

[12] Vera Hollink, Theodora Tsikrika, and Arjen de Vries. Semantic vs term-based query modification analysis. In *Proceedings of Tenth Dutch-Belgian Information Retrieval Workshop, Nijmegen, the Netherlands*, 2010.

[13] Vera Hollink, Theodora Tsikrika, and Arjen de Vries. The semantics of query modification. In *9th International Conference on Adaptivity, Personalization and Fusion of Heterogeneous Information (RIAO 2010)*, Paris, 28-30 April 2010. Accepted for publication.

[14] Jeff Huang and Efthimis N. Efthimiadis. Analyzing and evaluating query reformulation strategies in web search logs. In *Proceeding of the 18th ACM Conference on Information and Knowledge Management, Hong Kong, China*, pages 77–86, 2009.

[15] Michael Huggett and Joel Lanir. Static reformulation: a user study of static hypertext for query-based reformulation. In *Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital libraries, Vancouver, BC, Canada*, pages 319–328, 2007.

[16] Bernard J. Jansen. Search log analysis: what it is, what's been done, how to do it. *Library and Information Science Research*, 28(3):407–432, 2006.

[17] Bernard J. Jansen, Danielle L. Booth, and Amanda Spink. Patterns of query reformulation during web searching. *Journal of the Americal Society for Information Science and Technology*, 60(7):1358–1371, 2009.

[18] Bernard J. Jansen, Amanda Spink, and Jan O. Pedersen. The effect of specialized multimedia collections on web searching. *Journal of Web Engineering*, 3(3-4):182–199, 2004.

[19] Corinne Jörgensen and Peter Jörgensen. Image querying by image professionals. *Journal of the Americal Society for Information Science and Technology*, 56(12):1346–1359, 2005.

[20] Melanie Kellar, Kirstie Hawkey, Kori M. Inkpen, and Carolyn Watters. Challenges of capturing natural web-based user behaviours. *International Journal of Human Computer Interaction*, 24(4):385—409, 2008.

[21] Melanie Kellar, Carolyn Watters, and Michael Shepherd. A field study characterizing web-based information-seeking tasks. *Journal of the Americal Society for Information Science and Technology*, 58(7):999–1018, 2007.

[22] Tessa Lau and Eric Horvitz. Patterns of search: analyzing and modeling web query refinement. In *Proceedings of the Seventh International Conference on User Modeling, Banff, Canada*, pages 119–128, 1999.

[23] Julie B. Morrison, Peter Pirolli, and Stuart K. Card. A taxonomic analysis of what world wide web activities significantly impact people's decisions and actions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Seattle, WA, USA*, pages 163–164, 2001.

[24] Linking Open Data Project. Linked open data W3C SWEO community project, 2009. http://esw.w3.org/topic/SweoIG/TaskForces/CommunityProjects/LinkingOpenData, Last accessed February, 2010.

[25] Soo Young Rieh and Hong Xie. Analysis of multiple query reformulations on the web: The interactive information retrieval context. *Information Processing and Management*, 42(3):751–768, 2006.

[26] Guus Schreiber, Alia Amin, Lora Aroyo, Mark van Assem, Victor de Boer, Lynda Hardman, Michiel Hildebrand, Borys Omelayenko, Jacco van Osenbruggen, Anna Tordai, Jan Wielemaker, and Bob Wielinga. Semantic annotation and search of cultural-heritage collections: the MultimediaN E-Culture demonstrator. *Journal of Web Semantics*, 6(4):243–249, 2008.

[27] Abigail J. Sellen, Rachel Murphy, and Kate L. Shaw. How knowledge workers use the web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Minneapolis, MN, USA*, pages 227–234, 2002.

[28] Dian Tjondronegoro, Amanda Spink, and Bernard J. Jansen. A study and comparison of multimedia web searching: 1997-2006. *Journal of the American Society for Information Science and Technology*, 60(9):1756–1768, 2009.

[29] Martin Whittle, Barry Eaglestone, Nigel Ford, Valerie J. Gillet, and Andrew Madden. Data mining of search engine logs. *Journal of the American Society for Information Science and Technology*, 58(14):2382–2400, 2007.

[30] Hongyuan Zha, Xiaofeng He, Chris Ding, Horst Simon, and Ming Gu. Bipartite graph partitioning and data clustering. In *Proceedings of the Tenth International Conference on Information and Knowledge Management, Atlanta, GA*, pages 25–32, 2001.

10 Requested Budget

(1)

a)	appointment of research personnel	1 postdoc (3 years)
b)	additional travel budget	-
c)	project-related apparatus/software	-
d)	other related activities	-
	Total b, c, d	-

(2) Investments (incl. VAT)

-