In Proceedings of the International Workshop on Collaborative Knowledge Management for Web Information Systems (We.Know'07), held in conjunction with the 8th International Conference on Web Information Systems Engineering (WISE 2007)

How Do Users Express Goals on the Web? - An Exploration of Intentional Structures in Web Search

M. Strohmaier[1], M. Lux[2], M. Granitzer[3], P. Scheir[1,3], S. Liaskos[4], E. Yu[5]

[1]Graz University of Technology, 8010 Graz, Austria
[2]Klagenfurt University, 9020 Klagenfurt, Austria
[3]Know-Center Graz, 8010 Graz, Austria
[4]York University, Toronto, Canada
[5]University of Toronto, Toronto, Canada
markus.strohmaier@tugraz.at, mlux@itec.uni-klu.ac.at, mgrani@know-center.at, peter.scheir@tugraz.at, liaskos@yorku.ca, yu@fis.utoronto.ca

Abstract. Many activities on the web are driven by high-level goals of users, such as "plan a trip" or "buy some product". In this paper, we are interested in exploring the role and structure of users' goals in web search. We want gain insights into how users express goals, and how their goals can be represented in a semi-formal way. The paper presents results from an exploratory study that focused on analyzing selected search sessions from a search engine log. In a detailed example, we demonstrate how goal-oriented search can be represented and understood as a traversal of goal graphs. Finally, we provide some ideas on how to construct large-scale goal graphs in a semi-algorithmic, collaborative way. We conclude with a description of a series of challenges that we consider to be important for future research.
Keywords: information search, search process, goals, intentional structures

1 Motivation

In a highly influential article regarding the future of the web [1], Tim Berners-Lee sketches a scenario that describes a set of agents collaborating on the web to address different needs of users – such as "get medication", "find medical providers" or "coordinate appointments".

In fact, many activities on the web are already implicitly driven by goals today. Users utilize the web for buying products, planning trips, conducting business, doing research or seeking health advice. Many of these activities involve rather high-level goals of users, which are typically knowledge intensive and often benefit from social relations and collaboration. Yet, the web in its current form is largely non-intentional. That means the web lacks explicit intentional structures and representations, which would allow systems to, for example, associate users' goals with resources available on the web. As a consequence, every time users turn to the web for a specific purpose they are required to cognitively translate their high-level goals into the non-intentional structure of the web. They need to break down their goals into specific search queries, tag concepts, classification terms or ontological vocabulary. This prevents users from,

for example, effectively assessing the relevance and context of resources with respect to their goals, benefiting from the experiences of others who pursued similar goals and also prevents them from assessing conflicts or systematically exploring alternative means.
In a recent interview, Peter Norvig, Director of Google Research, acknowledged that understanding users' needs to a greater extent represents an "*outstanding*" research problem. He explains that Google is currently looking at *"finding ways to get the user more involved, to have them tell us more of what they want."* [2]. Having explicit intentional representations and structures available on the web would allow users to express and share their goals and would enable technologies and other users to explore, comprehend, reason about and act upon them.

It is only recently that researchers have developed a broad interest in the goals and motivations of web users. For example, several researchers studied intentionality and motivations in web search logs during the last years [3,4,5]. Because web search today represents a primary instrument through which users exercise their intent, search engines have a tremendous corpus of intentional artifacts at their disposal. We define intentional artifacts broadly to be electronic artifacts produced by users or user behaviour that contain *recognizable* "*traces of intent*", i.e. implicit traces of users' goals and intentions.

This paper represents our initial attempt towards exploring the role and structure of users' goals in web search queries. We want to learn in detail *how* users express their goals on the web - as opposed to *what* goals they have, which is in the focus of other studies [3,4,5]. We also want to explore how search goals can be represented in an explicit, semi-formal way and we are interested in learning about the different ways in which explicit goal representations could be useful, and to what extent. From our preliminary findings of an exploratory study, we want to give a qualitative account of identified potentials and obstacles in the context of goal-oriented search.

2 State of the Art

We will discuss two main streams of research that are relevant in the context of this paper: The first stream of research focuses on identifying and understanding *what* goals users pursue in web search. The second stream focuses on developing goal-oriented technical solutions, i.e. solutions that depend on the explicit articulation of user goals or automatic inference thereof.

In the first stream, researchers have proposed categories and taxonomies of user goals [4,5] and automatic classification techniques to classify search queries into goal categories [3]. Goal taxonomies include, for example, navigational, informational and transactional categories [3]. Different categories are assumed to have different implications on users' search behaviour and search algorithms. To give some examples: Navigational search queries (such as the query "citeseer") characterize situations where a user has a particular web site in mind and where he is primarily interested in visiting this page. Informational search queries (such as the query "increase wine crop") are queries where this is not the case, and users intend to visit multiple pages to, for example, learn about a topic [3]. Further research aims to empirically assess the distribution of different goal categories in search query logs via manual classification and subsequent statistical generalization [4] and/or Web Query Mining techniques [3,6]. There is some evidence that certain categories of goals can be identified algorithmically based on different features of user behaviour, such as "past user-click behaviour" and an analysis of "click distributions" [3]. Recently, a community of researchers with an interest in Query Log Analysis has formed at the World Wide Web 2007 conference as a separate workshop.

A second stream of research attempts to demonstrate the principle feasibility of implementing goal-orientation on an operational level. *GOOSE*, for example, is a prototypical goal-oriented search engine that aims to assist users in finding adequate search terms for their goals [7]. *Miro*, another example, is an application that facilitates goal-oriented web browsing [8]. The Lumiere Project focused on inferring goals of software users based on Bayesian user modeling [10]. Work on goal-oriented acquisition of requirements for hypermedia applications [11] shows that it is possible to translate high-level goals of stakeholders into (among other things) low level content requirements for web applications. Another example [12] facilitates purposeful navigation of geospatial data through goal-driven service invocation based on WSMO. WSMO is a web service description approach that decouples user desires from service descriptions by modeling low-level goals (such as "havingATripConfirmation") and non-functional property constructs [13]. In addition to these approaches, there have been several studies in the domain of information science that focus on different search strategies (such as top-down, bottom-up, mixed strategies) of users [14].

Apart from these isolated, yet encouraging, attempts, current research lacks a deep understanding about *how* users express their goals, and what explicit representations could be suitable to describe them.

3 How do Users Express Goals in Web Search?

We initiated an explorative study in response to the observation that there is a lack of research on *how* users express their goals in web search. In the following we will present preliminary findings from this study.

Data sources: We have used the AOL search database [15] as our main data source[1]. In addition to the AOL search database, several other web search logs are available [16]. We have used the AOL search database because it provides information about anonymous User IDs, time stamps, search queries, and clicked links. To our knowledge, the AOL search database is also the most recent corpus of search queries available (2006). We are aware of the ethic controversies arising from using the AOL search database. For example, although the User IDs are anonymous, a New York Times reporter was able to track back the identity of one of the users in the dataset [17]. As a consequence, we masked the search queries that are presented in

---

[1] Because the AOL search database was retracted from AOL shortly after releasing it, we obtained a copy from a secondary source: http://www.gregsadetsky.com/aol-data/ last accessed on July 15th, 2007.

this paper by maintaining their semantic frame structure, but exchanging certain frame element instantiations [19]. We will elaborate on this later on. In following such an approach, we aim to protect the real identity of the users being studied while retaining necessary temporal and intentional relations of search queries.

Methodology: In this study we were interested in how users express, refine, alter and reformulate their goals while searching. We have searched the AOL search database for different verbs that are considered to indicate the presence of goals, including verbs such as *achieve, make, improve, speedup, increase, satisfied, completed, allocated, maintain, keep, ensure* and others [18]. We subsequently annotated random results (different search queries) with semantic frame elements obtained from Berkeley's Framenet [19]. Framenet is a lexical database that aims to document the different semantic and syntactic combinatory possibilities of English words in each of its senses. It aims to achieve that by annotating large corpora of text. It currently provides information on more than 10.000 lexical units in more than 825 semantic frames [19]. A **lexical unit** is a pairing of a word with a meaning. For example, the verb "look" has several lexical units dealing with different meanings of this verb, such as "direct one's gaze in a specified direction" or "attempt to find". Each different meaning of the word belongs to a **semantic frame**, which is "a script-like conceptual structure that describes a particular type of situation, object or event along with its participants and props" [19]. Each of these elements of a semantic frame is called **frame elements**. Semantic frames are evoked by lexical units. To give an example, the semantic frame "Cause_change_of_position_on_a_scale" is evoked by a set of lexical units, such as decline, decrease, gain, plummet, rise, increase, etc, and has the core frame elements Agent [], Attribute [Variable], Cause [Cause] and Item [Item]. Agent refers to the person who causes a change of position on a scale, attribute refers to the scale that changes its value, cause refers to non-human causes to the change, and item refers to the entity that is being changed.

Example: The search query "Increase Computer Speed" can be annotated with Frame Elements from Framenet's lexicon. The lexical unit "increase" evokes the frame "Cause_change_of_position_on_a_scale", which we can use to annotate "Increase Computer Speed" in the following way: "**Increase** [$_{item}$ Computer] [$_{attribute}$ Speed]". The frame elements Agent and Cause do not apply here.

Selected Results: One verb we were using to explore the dataset was "increase". The query history depicted in Table 1 below presents an excerpt of the search history of a single user that performed search queries containing the verb "increase". We picked this particular search log because it demonstrates several interesting aspects of the role of goals in web search. We do not claim that this user's search behaviour is typically or representative for a larger set of users or queries. In fact, the majority of search queries in the AOL search database is of a non-intentional nature. We discuss the implications of this observation in the Section 5.

We obtained the complete search record of the selected user, frame-annotated his intentional queries based on the FrameNet lexicon and classified the queries from an intentional perspective (e.g. refinement, generalization, etc). The particular frame used during annotation was "Cause_change_of_position_on_a_scale", which is evoked by the verb "increase". For privacy reasons, we modified the search queries in the following way: We retained the verbs and attributes which were part of the original query, but modified the contents of the semantic frame element item (e.g. wine crop) and cause (e.g. fertilizer) as well as time stamps (maintaining relative time differences with an accuracy of +/- 60 seconds). We'd like to remark that the users' search history below was interrupted by other, non-intentional queries (queries such as "flickr.com") and also other more complex intentional queries. For reasons of illustration and simplicity, we leave these out in Table 1.

Nr.	Query	Frame Annotation	Time Stamp	Goal
#1	How to get more wine crop	How to **get more** [$_{item}$wine crop]	2006-03-30 19:29:59	Formulation
#2	Fertilizer or insecticide to increase wine crop	[$_{cause}$ Fertilizer] or [$_{cause}$ insecticide] to **increase** [$_{item}$wine crop]	2006-03-30 19:45:28	Refinement
#3	Fertilizer to increase wine crop	[$_{cause}$ Fertilizer] to **increase** [$_{item}$ wine crop]	2006-03-30 19:46:11	Refinement
[further non-intentional queries, not related to wine crop]				
#4	Increase wine crop	**increase** [$_{item}$ wine crop]	2006-03-30 19:48:25	Generali-zation
#5	How to get rich wine crop	How to **get rich** [$_{item}$ wine crop]	2006-04-07 06:29:19	Different Goal Formulation
[non-intentional query "wine crop"]				
#6	How to get good wine crop	How to **have good** [$_{item}$ wine crop]	2006-04-07 06:40:45	Re-formulation
[further non-intentional queries and further more complex intentional queries related to "wine crops"]				

Table 1. Frame-based Annotation of Selected Queries from a Single Search Session

From a semantic frame perspective, it is interesting to see that it is not possible to annotate all of the above queries consistently. While the verb **increase** evokes the corresponding frame "Cause_change_of_position_on_a_scale" in queries #2, #3 and #4, the other queries #1, #5, and #6 do not contain **increase** and therefore do not evoke the same frame. Although FrameNet contains lexical entries for the verbs **get** and **have** and the adjectives **good**, **rich** and **more**, the word senses *get more*, *get rich* and *have good* are not yet captured as lexical units in the FrameNet lexicon. However, it is easily conceivable that an expanded or customized version of FrameNet (possibly in combination with WordNet) would contain these units and that they could be associated with the same semantic frame.

From a goal-oriented perspective, we will use our findings to develop a set of hypothesis that we believe are relevant and helpful to further study the role and structure of users' goals on the web.

Several things are noteworthy in the search history of the above user: First, the user started off with a goal formulation (#1 how to get more wine crop) and then proceeded with a refinement of this goal in a second query (#2 Fertilizer or insecticide to increase wine crop). The provided time stamps reveal that in this case, the time difference between the two queries was more than 15 minutes! Although it is hard to assess the real cause for this time lag, the AOL search database provides a possible explanation by listing the websites that the user visited in response to query #1, which includes a discussion board website hosting discussions on different strategies to get more "wine crop" (including "insecticides" and "fertilizer"). This allows us to hypothesize that **H1: Goal refinement is a time-intensive process during search**.

In query #3, the user performed a further refinement of his goal to "fertilizer to increase wine crop" and in #4, he performs a generalization to "Increase wine crop". This is interesting again from a goal-oriented perspective: Instead of refining his goals in a strict top-down approach, the user alternates between top-down (refining) and bottom up (generalizing) goal formulations. We consider this observation in a hypothesis 2 that claims that, from a goal-oriented perspective, user search is neither a strict top-down, nor a purely bottom-up approach, but a combination of both. While we focus on informational queries only, previous studies have found that the type of approach does not only depend on the type of task, but also different types of users [14]. This leads us to hypothesize **H2: Users search by iteratively refining, generalizing and reformulating goals, in no particular order**.

In query #5 the user performs a different goal formulation: "How to get rich wine crop". Instead of focusing on quantity ("get more" / "increase"), the search now can be interpreted to focus on the quality of wine crop ("get rich"). In query #6, a goal re-formulation is performed. This can be regarded to represent the same goal, but articulated in a slightly different way ("get good" instead of "get rich" wine crop). Another very interesting observation is that there is a time span of more than 7 days between queries #1-#4 and queries #5-#6! Although we have no information about what the user might have done in between these search activities, we use this evidence to tentatively hypothesize that identifying different, but related, goals is difficult for users, and it involves significant time and potentially cognitive efforts. In a more intuitive way, we can say that it seems that, especially with high-level, knowledge intensive goals, users learn about their goals as they go. We formulate this observation in hypothesis **H3: Exploring related goals is more time-intensive than goal refinement**.

And finally, we can observe that a smaller amount of time is passing between search queries #5 and #6. The question that is interesting to ask based on this observation is whether goal refinements require more time and cognitive investments from users than goal re-formulations. One might expect that users with search experience become skilled in tweaking their queries based on the search engines' responses without modifying their initial goal. We express this question in our hypothesis **H4: Goal re-formulation requires less time than goal exploration or goal refinement**. Next, we will explore some implications of these observations.

Analysis: If hypothesis H1 would be corroborated in future studies, offering users possible goal refinements would be very likely to be considered a useful concept. If hypothesis H2 would be supported in further studies, goal-oriented search would not only need to focus on goal refinement, but also on providing a range of different intentional navigation structures, allowing to flexibly alternate between refining, generalizing and exploring goals. If the exploration of goals represents a very time intensive process (H3), then users can be assumed to greatly benefit from having access to the goals of other users. And finally, if goal re-formulation does not require significant amounts of time (H4), there might be little motivation for researchers to invest in semantic similarity of web searchers, but more motivation to invest in intentional similarity.

Surprisingly, when analyzing current search technologies such as Google, we can see that there is almost no support for any of these different goal-related search tasks (refinement, generalization, etc) identified. Although Google helps in reformulating search queries ("Did you mean X?"), this – at most – can be regarded to provide some support for users in goal re-formulation on a *syntactic* level, but not on a truly *intentional* level (help in goal refinement, generalization, etc).

These observations immediately raise a set of interesting research questions: Do the formulated hypotheses hold for large sets of search sessions? How can the hypotheses be further refined to make them amenable to algorithmic analysis? And how can the identified goals be represented in more formal structures? While we are interested in all of these questions, in this paper we will only discuss the issue of more formal representations in some greater detail.

4 Representing Search Goals as Semi-Formal Goal Graphs

We have modeled the goals of a user who is interested in "wine crop" with the agent-and goal-oriented modeling framework i* [20]. When applying i*, we focused on goal aspects and neglected agent-related concepts such as actors, roles and others. The i* framework provides elements such as softgoals, goals, tasks, resources and a set of semantic relations between them. The goal graph in Figure 1 was constructed by one of the authors of this paper based on the frame-annotated goals depicted in Table 1. In the diagram, the goals of the users are represented through oval-shaped elements. Means-ends links are used to indicate alternative ways (means) by which a goal (ends) can be fulfilled. Goals represent states of affairs to be reached, and tasks, which are represented through hexagonal elements, describe specific activities that can be performed for the fulfillment of goals. Soft-goals, which are represented through cloud-shaped elements, describe goals for which there is no clear-cut criterion to be used for deciding whether they are satisfied or not. Thus, soft-goals are fulfilled or denied to a certain degree, based on the presence or absence of relevant evidence. In i* diagrams, links such as "help" or "hurt" are used to represent how a belief about the fulfillment or denial of a soft-goal depends on the satisfaction of other goals. From the goal-graph in Figure 1 we can infer that the goal "increase wine crop" can be achieved through a variety of means: Fertilizer, Insecticides and Irrigation all represent means to achieve the end of increasing wine crop. The goal "Increase wine crop" and the related goal "Improve wine crop" both have "help" contribution links to the overarching soft-goal "Winery be successful".

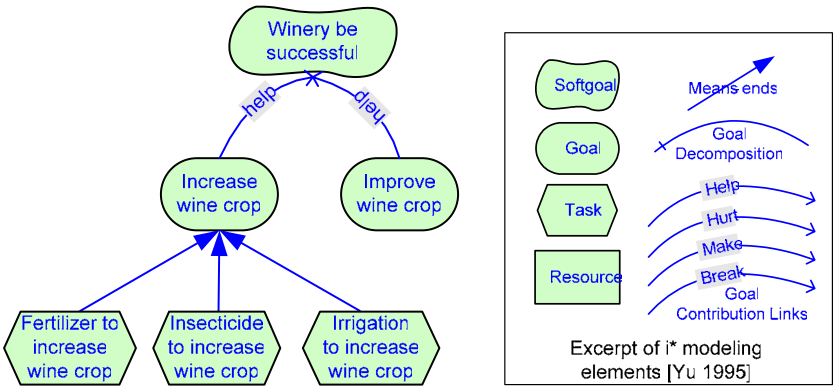

Fig. 1. Representing Users' Search Goals in a Semi-Formal Goal Graph

Assuming that such goal graphs can be constructed for a range of different domains (which is evident in a broad set of published examples from the domain of requirements engineering), it would be interesting to see how the different goal-related activities of users during search (such as goal formulation, goal refinement, goal generalization, etc) can be represented as a traversal of such a goal graph. We will explore this question next.

4.1 How Can Search be Understood as a Traversal through A Goal Graph?

Modifying search engines' algorithms to exploit knowledge about users' goals has a high priority for search engine vendors [5]. Being able to relate search queries to nodes in a goal graph could enable search engines to provide users goal-oriented support in search. This could mean that software could offer users to refine their search goals, generalize them or propose related goals from other users.

Figure 2, depicts the results of manually associating the search queries presented in Table 1 with the goal graph introduced in Figure 1. We can see that the user starts his search by formulating a version of the goal "increase wine crop" in query #1. This goal is refined in query #2 "Fertilizer or insecticides to increase wine crop" which can be mapped onto the two means "Fertilizer to increase wine crop" and "Insecticides to increase wine crop". Query #3 "fertilizer to increase wine crop" represents a further refinement. In query #4, the user generalizes his search goal to "increase wine crop" again. Query #5 and #6 relate to a different goal: "Improve wine crop". Query #5 and #6 can be considered to be re-formulations of the same goal.

Interestingly, the goal graph reveals that the user did not execute search queries related to the means "Irrigation to increase wine crop" or the soft-goal "Winery be successful", although one can reasonably expect that the user might have had a genuine interest in these goals too (although validation of this claim is certainly hard without user interaction).

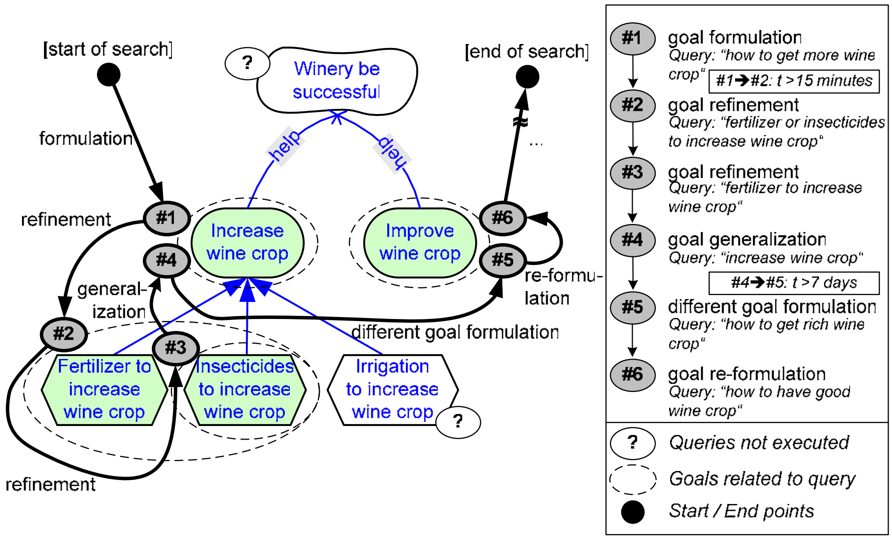

Fig. 2. Goal-Oriented Search as a Traversal of Goal Graphs

As a consequence, a major benefit of having goal graphs available during search could be pointing users to refined goals or making sure that users do not miss related goals. But assuming that having such goal graphs would be beneficial, how can they be constructed?

4.2 How Can Large-Scale Goal Graphs be Constructed?

Mapping search queries onto goal graphs presumes the existence and availability of goal graphs. In our example, we have hand-crafted a goal graph for illustration purposes. However, manually constructing such goal graphs is costly, and anticipating the entirety, or even a large proportion, of users' goals on the internet would render such an approach unfeasible. So how can we construct large-scale goal graphs that do not rely on the involvement of expert modelers? Automatic user goal identification is an open research problem [6], and answering this question satisfyingly would go well beyond the scope of this paper, but we'd like to discuss some pointers and ideas: The recent notion of folksonomies has powerfully demonstrated that meaningful relations can emerge out of collective behaviour and interactions [21]. We would like to briefly explore this idea and some of its implications for constructing large-scale goal graphs based on frame-analysis of intentional artifacts.

Let's assume that a system has the capability to come up with frame-based annotations of search queries. The search query "fertilizer or insecticide to increase wine crop" would then be annotated in a way that is depicted on the left side of Figure 3. Based on such annotations, a goal graph construction algorithm could use heuristics to construct a goal graph similar to the one depicted on the right side of Figure 3.

Heuristic rules could, for example, prescribe that the root goal is represented by the central verb ("increase") and its corresponding item ("wine crop"), and that the means to this end are represented by the frame elements cause ("fertilizer", "insecticide"). Each time a user formulates an intentional search query, the goal graph construction algorithm could construct such small, atomic goal graphs heuristically.

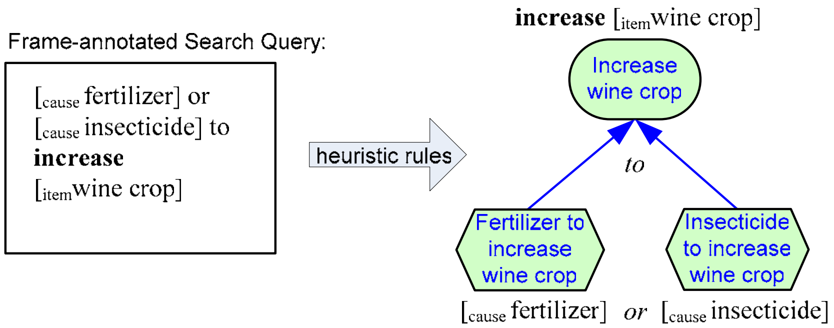

Fig. 3. Heuristic Construction of Atomic Goal Graphs via Frame-Annotation of Search Queries

In a next step, these atomic goal graphs constructed from different users' search queries would need to be connected to larger whole. Considering hypothesis 2, this appears to be a task that is hard to perform by algorithms alone. Nevertheless, usage data analysis, explicit user involvement or semi-automatic, collaborative model construction efforts (as e.g. pursued by the ConceptNet project [9]) might help to overcome this issue, which can be considered to represent a non-trivial research challenge.

5 Implications and Threats to Validity

We are aware that our particular research approach puts some constraints on the results of our work: Due to our focus, the search queries we analyzed were not required to be representative and, in fact, they are not. To obtain some quantitative evidence, two of the authors have categorized a pseudo-random sample (based on java.util.Random randomizer) of 2000 out of 21,011,340 queries into intentional and non-intentional categories, based on the criterion whether a query contains *at least* one verb (infinitive form, excluding gerund) and *at least* one noun. For each of these candidates, two authors of this paper judged whether it would be possible to envisage the goal a user might have had based on a specific query (such as "increase computer speed"). From our analysis, only 2.35% (47 out of 2000) of the searches from the AOL search database can be considered to be such "intentional queries". The probability of occurrence then results in a 95% confidence interval of [0.0169, 0.0301] for the probability of a query being intentional according to our criteria.
In contrast to these findings, related studies found somewhat higher numbers. A study reported in [4] suggests that 35% of search sessions have a general, high-level information research goal (such as questions, undirected requests for information, and advice seeking). The difference in numbers might be explained by different levels of analysis and a more relaxed understanding of goals in [4], which allows a broader set of queries (including queries that do not have verbs) to be labelled as goal-related.

There are several implications of this discrepancy: While users often have high-level goals when they are searching the web, they are currently not rewarded for formulating (strictly) intentional queries. In fact, one can assume that formulating non-intentional queries represents a (locally) successful strategy in today's search engine landscape. As a result, users might have adapted to the non-intentional mode in which Google, Yahoo and other search engines operate today. However, this situation makes it necessary for users to cognitively translate their high-level goals into search queries and perform reasoning about their goals in their mind. This potentially increases the cognitive burden of users and makes it hard for systems to connect them with other users who pursue similar goals or allowing them to benefit from the experiences made by other searchers.

We do not believe that these implications put constraints on our results: With a collaborative goal modeling approach, even a small percentage of strictly intentional queries could be used to construct large-scale goal graphs. Even if the percentage of intentional queries among the entirety of search queries would be as low as 1% or even lower, the sheer amount of queries executed on the World Wide Web would still provide algorithms with a rich corpus to construct large-scale goal graphs. On the web, such an approach is by far not unusual: For example, on wikipedia, a minority of users contributes content that is being used by a majority. However, the task of constructing large-scale goal graphs would obviously become much easier if users actually would be aware that search engines would interpret their queries as an expression of intent rather than an input that is being used for text string matching.

6 Conclusions

Based on our preliminary findings, we can formulate a set of interesting research challenges: First, how can large-scale goal graphs be represented and constructed? How can intentional artifacts (such as search queries) be associated with nodes in such goal graphs? How can goals and web resources be associated? And how can collaboration on the internet support the construction of such intentional structures?

Our work represents an initial attempt towards understanding the role and structure of goals in web search. We have demonstrated how search processes can be understood as a traversal through goal graphs and have provided some ideas on how to construct large scale goal graphs. In future work, we are interested in further investigating and shaping intentional structures on the web.

References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. Scientific American 284 (2001)
2. Greene, K., The Future of Search. http://www.technologyreview.com/Biztech/19050/, last accessed on July 18th, 2007, MIT Technology Review, July 16 (2007)
3. Lee, U., Liu, Z., Cho, J.: Automatic Identification of User Goals in Web Search. In: WWW '05: Proceedings of the 14th International World Wide Web Conference, New York, NY, USA, ACM Press (2005) 391–400

4. Rose, D., Levinson, D.: Understanding User Goals in Web Search. In Feldman, S.I., Uretsky, M., Najork, M., Wills, C., eds.: Proceedings of the 13th International World Wide Web Conference, ACM Press (2004) 13–19
5. Broder, A.: A Taxonomy of Web Search. SIGIR Forum 36 (2002) 3–1
6. Baeza-Yates, R., Calderon-Benavides, L., Gonzalez-Caro, C.: The Intention Behind Web Queries. In Crestani, F., Ferragina, P., Sanderson, M., eds.: Proceedings of String Processing and Information Retrieval (SPIRE). Volume 4209 of Lecture Notes in Computer Science., Springer (2006) 98–109
7. Liu, H., Lieberman, H., Selker, T.: GOOSE: A Goal-Oriented Search Engine with Commonsense. In: AH ’02: Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, London, UK, Springer-Verlag (2002) 253–263
8. Faaborg, A., Lieberman, H.: A Goal-Oriented Web Browser. In: CHI ’06: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, New York, NY, USA, ACM Press (2006) 751–760
9. Liu, H., Singh, P.: Conceptnet - A Practical Commonsense Reasoning Tool-Kit. BT Technology Journal 22 (2004) 211–226
10. Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, K.: The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. In: In Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Madison, WI (1998) 256–265
11. Bolchini, D., Paolini, P., Randazzo, G.: Adding Hypermedia Requirements to Goal-Driven Analysis. In: In Proceedings of the 11th IEEE International Conference on Requirements Engineering (RE 2003), IEEE Computer Society (2003) 127–137
12. Tanasescu, V., Gugliotta, A., Domingue, J., Villarıas, L., Davies, R., Rowlatt, M., Richardson, M., Stincic, S. In: Geospatial Data Integration with Semantic Web Services: the eMerges Approach. (2007)
13. Roman, D., Keller, U., Lausen, H., de Bruijn, J., Lara, R., Stollberg, M., Polleres, A., Feier, C., Bussler, C., Fensel, D.: Web Service Modeling Ontology. Applied Ontology 1 (2005) 77–106
14. Navarro-Prieto, R., Scaife, M., Rogers, Y.: Cognitive Strategies in Web Searching. In: Proceedings of the 5th Conference on Human Factors & the Web. (1999)
15. Pass, G., Chowdhury, A., Torgeson, C.: A picture of search. Proceedings of the 1st International Conference on Scalable Information Systems, ACM Press New York, NY, USA, (2006)
16. Jansen, B., Spink, A.: How Are We Searching the World Wide Web? A Comparison of Nine Search Engine Transaction Logs. Information Processing and Management 42 (2006) 248–263
17. Barbaro, M., Zeller Jr, T.: A Face Is Exposed for AOL Searcher No. 4417749, New York Times, August 9 (2006)
18. Regev, G., Wegmann, A.: Where Do Goals Come From: The Underlying Principles of Goal-Oriented Requirements Engineering. In: RE ’05: Proceedings of the 13th IEEE International Conference on Requirements Engineering (RE’05), Washington, DC, USA, IEEE Computer Society (2005) 253–362
19. Ruppenhofer, J., Ellsworth, M., Petruck, M., Johnson, C., Scheffczyk, J.: FrameNet II: Extended Theory and Practice, International Computer Science Institute, University of California at Berkeley (2006)
20. Yu, E.: Modelling Strategic Relationships for Process Reengineering. PhD thesis, Department of Computer Science, University of Toronto, Toronto, Canada (1995)
21. Mika, P.: Ontologies Are Us: A Unified Model of Social Networks and Semantics. In: International Semantic Web Conference. LNCS, Springer (2005) 522–536