

Understanding User Goals in Web Search

Daniel E. Rose

Yahoo! Inc.

701 First Avenue, MS B201

Sunnyvale, CA 94089 USA

+1 408 349 7992

drose@yahoo-inc.com

Danny Levinson

Yahoo! Inc.

144 Fourth Avenue SW, Suite 2600

Calgary AB T2P 3N4 Canada

+1 403 303 4590

dlevinso@yahoo-inc.com

ABSTRACT

Previous work on understanding user web search behavior has focused on how people search and what they are searching for, but not why they are searching. In this paper, we describe a framework for understanding the underlying goals of user searches, and our experience in using the framework to manually classify queries from a web search engine. Our analysis suggests that so-called “navigational” searches are less prevalent than generally believed, while a previously unexplored “resource-seeking” goal may account for a large fraction of web searches. We also illustrate how this knowledge of user search goals might be used to improve future web search engines.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – *search process*; H.4.m [Information Systems Applications]: Miscellaneous.

General Terms

Measurement, Experimentation, Human Factors

Keywords

Web search, information retrieval, user behavior, user goals, query classification.

1. INTRODUCTION

If we imagine seeing the world from the perspective of a search engine, our only view of user behavior would be the stream of queries users produce. Search engine designers often adopt this perspective, studying these query streams and trying to optimize the engines based on such factors as the length of a typical query. Yet this same perspective has prevented us from looking beyond the query, at *why* the users are performing their searches in the first place.

The “why” of user search behavior is actually essential to satisfying the user’s information need. After all, users don’t sit down at their computer and say to themselves, “I think I’ll do some searches.” Searching is merely a means to an end – a way to satisfy an underlying goal that the user is trying to achieve. (By “underlying goal,” we mean how the user might answer the question “why are you performing that search?”) That goal may be choosing a suitable wedding present for a friend, learning which local colleges offer adult education courses in pottery, seeing if a favorite author’s new book has been released, or any

number of other possibilities. In fact, in some cases the same query might be used to convey different goals – for example, the query “ceramics” might have been used in any of the three situations above (assuming it is also the title of the book in question).

What difference would it make if the search engine knew the user’s goal? At the very least, the engine might provide a user experience tailored toward that goal. For example, the display of relevant advertising might be welcome in a shopping context, but unwelcome in a research context. In fact, we have argued elsewhere [10] that goal-sensitivity will be one of the crucial factors in future search user interfaces. But the potential to capitalize on this goal sensitivity goes beyond the user interface. The underlying relevance-ranking algorithms that determine which results are presented to users might differ depending on the search goal. For example, queries that express a need for advice might rely more on usage- or connectivity-based relevance factors, while those involving open-ended research might weight traditional information retrieval measures (such as term frequency) more highly.

Our premise is that web searches reflect a diverse set of underlying user goals, and that knowledge of those goals offers the prospect of future improvements to web search engines. Achieving these improvements is an ambitious project involving three primary tasks. First, we need to create a conceptual framework for user goals. Second, we need a way for search engines to associate user goals with queries. Third, we need to modify the engines in order to exploit the goal information.

In this paper we focus on the first task, and the initial parts of the second: characterizing user search goals and examining the problem of inferring goals from query behavior. We begin in section 2 by looking at previous work on understanding information-seeking behavior. Next, in section 3, we describe our model of search goals. In section 4, we review the methodology used to classify queries using our model, and we provide some results from this analysis. We conclude with some final thoughts about the applicability of this work.

2. RELATED WORK

Studies of user search behavior have a long history in Information and Library Science. These include studies of the reference interview process, long before most users had access to computer-assisted search tools. When search engines first became available for use by researchers, many studies were conducted that attempted to understand user search behavior in an online context. For example, Bates [4] looked at the different ways in which people performed searches, and later proposed ways to characterize the overall search process [5]. Belkin’s Anomalous States of Knowledge (ASK) framework was an early attempt to

Copyright is held by the author/owner(s).

WWW 2004, May 17–22, 2004, New York, New York, USA.

ACM 1-58113-844-X/04/0005.

model the cognitive state of the user and then translate this understanding into a practical design for an information retrieval system [6]. Included in the ASK study was an analysis of some of the different types of information needs of different users. For example, one type of ASK was summarized as “Well-defined topic and problem,” while another was “Information needed to produce directions for research.”

Once web search engines became available and popular, studies of web search behavior followed quickly. For example, Silverstein et al. conducted an analysis of query logs from the AltaVista search engine, confirming some of the original findings of web search use, such as the predominance of very short queries [11]. A summary of many of the early studies may be found in Jansen and Pooch’s 2000 review [9].

One of the most comprehensive attempts to understand web search behavior has been the ongoing research of Spink and her colleagues, who analyzed query logs of the Excite search engine from 1997, 1999, and 2001 [13]. Although there have been some changes in user behavior during this period (such as a decrease in willingness to look at more than one page of search results), Spink et al. found that general search strategies have remained fairly constant.

Prior to the advent of the worldwide web, search engine designers could safely assume that users had an “informational” goal in mind. That is, users’ reason for searching was generally to “find out about” their search topic. This was due both to the nature of the population with access to full-text search engines (students, researchers, lawyers, intelligence analysts, etc.) and to the nature of the databases that could be searched (with services such as Westlaw, Dialog, Medline, Lexis/Nexis, etc.)

But in the web era, search engines are used for more than just research. Even the most cursory look at the query logs of any major search engine makes it clear that the goals underlying web searches are many and varied. And while the vast body of work described above has helped us to understand *what* users are searching for and *how* their information-seeking process works, there have been few attempts to look at *why* users are searching.

One of the few exceptions is Broder’s “Taxonomy of Web Search” [7]. Motivated by the idea that the traditional notion of an “information need” might not adequately describe web searching, Broder came up with a trichotomy of web search “types”: navigational, informational, and transactional. *Navigational* searches are those which are intended to find a specific web site that the user has in mind; *informational* searches are intended to find information about a topic; *transactional* searches are intended to “perform some web-mediated activity.”

3. A FRAMEWORK FOR SEARCH GOALS

Our first task was to understand the space of user goals. In particular, we needed to come up with a framework that could identify and organize a manageable set of canonical goal categories. These goal categories, in turn, must encompass the majority of actual goals users have in mind when searching.

To develop the goal framework, we looked at a sample of queries from the AltaVista search engine [1]. We brainstormed a variety of goal possibilities, based on our own experiences, some previous internal query analysis at AltaVista, and a preliminary examination of the query set. This resulted in a flat list of goals.

This list served as a basis for an initial goal classification framework, which we then used to categorize a sample of 100-200 queries. Next, we revised the framework to accommodate the results of the classification test. Categories were modified, or new categories added, when queries did not fit the existing framework. Some goal categories proposed early on, such as “finding a place in the world” (e.g. a map request), were dropped as unrepresentative. Some categories were merged, some were split more finely, and some entirely new ones arose. This propose-classify-refine cycle was repeated three times, each with a new set of queries.

One of our early findings was that there were many cases where the goal of the search was neither to find a web site nor to get information, but simply to get access to an online resource. For example, a query such as **beatles lyrics** suggests not a desire to learn *about* lyrics to Beatles songs, but simply a desire to view the lyrics themselves. This led to the creation of a broad new goal category that we call *resource* searches. We believe these resource searches are a relatively neglected category in the search engine world.

As we repeatedly revised the set of goal categories, we gradually reached the conclusion that the goals naturally fell into a hierarchical structure. In fact, the top level of the hierarchy resembles Broder’s trichotomy, but our more general “resource” category replaces his notion of “transactional” queries. Our resulting goal framework is shown in Table 1.

We define the **navigational** goal as demonstrating a desire by the user to be taken to the home page of the institution or organization in question. To be considered navigational, the query must have a single authoritative web site that the user already has in mind. For this reason, most queries consisting of names of companies, universities, or well-known organizations are considered navigational. Also for this reason, most queries for people – including celebrities – are not. A search for celebrities such as Cameron Diaz or Ben Affleck typically results in a variety of fan sites, media sites, and so on; it’s unlikely that a user entering the celebrity name as a query had the goal of visiting a specific site.

Informational queries are all focused on the user goal of obtaining information about the query topic. This category includes goals for answering questions (both **open-** and **closed-ended**) that the user has in mind, requests for **advice**, and **“undirected”** requests to simply learn more about a topic. Undirected queries may be viewed as requests to “find out about” or “tell me about” a topic; most queries consisting of topics in science, medicine, history, or news qualify as undirected, as do the celebrity queries mentioned above. Note that the two question-goal categories do not require that the user explicitly express the query in the form of a question; the query “last czar of russia” is reasonably interpreted as a closed-class question “who was the last czar of Russia?” Similarly, queries in the “advice” category may take many forms.

The informational goal class also includes the desire to **locate** something in the real world, or simply get a **list** of suggestions for further research. Most product or shopping queries have the “locate” goal – I’m searching the web for X because I want to know where I can buy X. Plural query terms are a highly reliable indicator of the list goal.

**Table 1: The Search Goal Hierarchy. Queries are only assigned to leaf nodes.
All examples are taken from actual AltaVista queries.**

SEARCH GOAL	DESCRIPTION	EXAMPLES
1. Navigational	My goal is to go to specific known website that I already have in mind. The only reason I'm searching is that it's more convenient than typing the URL, or perhaps I don't know the URL.	aloha airlines duke university hospital kelly blue book
2. Informational	My goal is to learn something by reading or viewing web pages	
2.1 Directed	I want to learn something in particular about my topic	
2.1.1 Closed	I want to get an answer to a question that has a single, unambiguous answer.	what is a supercharger 2004 election dates
2.1.2 Open	I want to get an answer to an open-ended question, or one with unconstrained depth.	baseball death and injury why are metals shiny
2.2 Undirected	I want to learn anything/everything about my topic. A query for topic X might be interpreted as "tell me about X."	color blindness jfk jr
2.3 Advice	I want to get advice, ideas, suggestions, or instructions.	help quitting smoking walking with weights
2.4 Locate	My goal is to find out whether/where some real world service or product can be obtained	pella windows phone card
2.5 List	My goal is to get a list of plausible suggested web sites (I.e. the search result list itself), each of which might be candidates for helping me achieve some underlying, unspecified goal	travel amsterdam universities florida newspapers
3. Resource	My goal is to obtain a resource (not information) available on web pages	
3.1 Download	My goal is to download a resource that must be on my computer or other device to be useful	kazaa lite mame roms
3.2 Entertainment	My goal is to be entertained simply by viewing items available on the result page	xxx porno movie free live camera in l.a.
3.3 Interact	My goal is to interact with a resource using another program/service available on the web site I find	weather measure converter
3.4 Obtain	My goal is to obtain a resource that does not require a computer to use. I may print it out, but I can also just look at it on the screen. I'm not obtaining it to learn some information, but because I want to use the resource itself.	free jack o lantern patterns ellis island lesson plans house document no. 587

Resource queries all represent a goal of obtaining something (other than information). If the resource is something I plan to use in the offline world, such as song lyrics, recipes, sewing patterns, etc., we call it an “**obtain**” goal. If the resource is something that needs to be installed on my computer or other electronic device to be useful, the goal is “**download**.” If my goal is simply to experience (typically view or read) the resource for my enjoyment, the goal is “**entertain**.” The most common example of queries with an entertainment goal were those that dealt with pornography. Finally, the “**interact**” goal occurs when the intended result of the search is a dynamic web service (such as a stock quote server or a map service) that requires further interaction to achieve the user’s task.

The search goal framework described above proved to be both stable (requiring no major revisions as new queries were encountered) and comprehensive (encompassing the goals of all the queries we had seen). We were therefore able to move on to the second major task, associating goals with queries.

4. ASSOCIATING GOALS WITH QUERIES

There are two ways a search engine might associate goals with queries at runtime: either the user can identify the goal explicitly through the user interface, or the system can attempt to infer the goal automatically. Google’s “I’m feeling lucky” feature [8], in which users implicitly identify their goal as “navigate to a specific web site,” may be thought of as an early example of the first

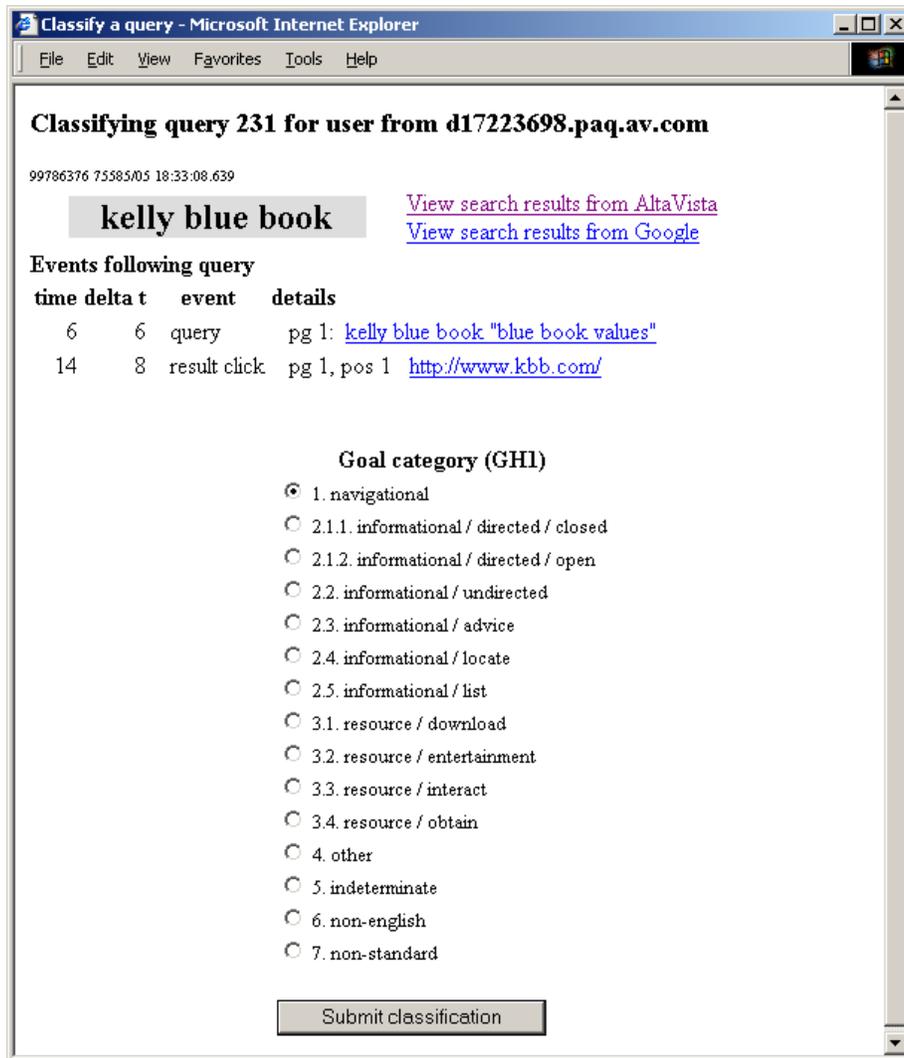


Figure 1: A screenshot of the tool used to assist manual query classification.

approach. The second approach would involve automatic classification using statistical or machine learning methods; these methods in turn will require hundreds or thousands of examples of classified queries (and their associated features) as training examples.

In either case, we need to know the relative prevalence of various goals. And if we hope to infer goals automatically in the future, we need to know that it is possible to do so manually. This section describes our work on these initial aspects of the problem; the remaining parts of the task will be the focus of future work.

4.1 Manual Query Classification

In order to definitively know the underlying goal of every user query, we would need to be able to ask the user about his or her intentions. Clearly, this is not feasible in most cases. But can the goal be determined simply by looking at the query itself, or is more information required?

We believe that in many cases, user goals can be deduced from looking at user behavior available to the search engine. Included in this behavior are the following:

- the query itself
- the results returned by the search engine
- results clicked on by the user
- further searches or other actions by the user.

We wanted to determine whether this was sufficient information for a human to consistently classify queries according to our goal framework. Once we could successfully classify queries manually, we would be able to provide training data for a future automatic classification system.

To facilitate the task of manual classification, we created a research tool that provided these four types of information for sets of queries taken from the AltaVista query logs. A screen shot of the classification tool interface is shown in Figure 1.

Table 2: Events following the query *final fantasy*.

Time	Delta t	Event	Details
36	36	result click	pg 1, pos 1 http://www.ffonline.com
113	77	query	pg 1 <i>final fantasy</i>
118	5	result click	pg 1, pos 8 http://www.eyesonff.com
147	29	result click	pg 1, pos 8 http://www.eyesonff.com

The query (**kelly blue book** in this example) appears in the gray-highlighted box at upper left. To the right of the query are links which lead to the search results that appear when the query is executed on two major search engines. Beneath this is a table of search engine events (clicks and queries) that this user performed following the initial query. In this case, we see that six seconds after issuing the query, the user entered a new, more specific query on the same topic. (The syntax suggests that this query resulted from the user clicking on a suggested query refinement term using AltaVista’s Prisma [2, 3] assisted search tool.) Eight seconds later, he or she clicked on the first result, **www.kbb.com**, which is the home page of the Kelly Blue Book (a publication that gives guidelines for new and used car prices). Thus a human classifier using the tool (namely, one of the authors) concluded that the underlying user goal for this query was “navigational,” and selected the corresponding radio button. When the “Submit classification” button is pressed, a new query is displayed, together with its corresponding information. In the example shown, a human classifier could probably have guessed the goal simply by viewing the initial query. Yet there are cases where each of the sources of information played a role in assessing the user’s goal.

Consider the query **final fantasy**. This is the name of a series of popular computer games. Did the user want to find a place to buy one of the games (a “locate” goal)? Did he or she intend to go to some official Final Fantasy web site (a “navigational” goal)? A look at the search results on AltaVista and Google shows that there isn’t an authoritative web site for the game. The game’s manufacturer has a web site, but it covers many games, has no specific page for the entire Final Fantasy series, and is ranked #3 on both AltaVista and Google. This casts some doubt on likelihood of a navigational goal. The result list contains many sites with information about the games, and many sites where one can buy the games. The user’s event history, shown in Table 2, provides further information.

The user examined the result list for 36 seconds, then visited the web site **www.ffonline.com**, described as “an unofficial guide to Final Fantasy.” About a minute later, s/he returned to the original query, and then chose a different web site, “Eyes on Final Fantasy,” (**www.eyesonff.com**), a site containing news and information about the games. This pattern indicates that the user was not interested in buying the game, but simply wanted some sort of information about it – perhaps the latest news about future releases. In this case, we’d conclude that the underlying goal was “undirected” information.

4.2 Results

Three sets of approximately 500 U.S. English queries¹ each were randomly selected from the AltaVista query logs on different days and at different times of the day. These were manually classified, one set using the classification tool as described above, and two sets using an earlier version that did not contain the user’s event history. Results are shown in Table 3. (Note that the “open” and “closed” categories have been collapsed into a single “directed” category, due to the low numbers of results.)

It is interesting to note that nearly 40% of queries were non-informational in every case, and a large fraction of the informational queries appeared to be attempts to locate a product or service rather than to learn about it. In fact, just over 35% of all queries appeared to have the kind of general research goal (questions, undirected requests for information, and advice-seeking) for which traditional information retrieval systems were designed.

It is also interesting that the relative distributions of goal categories are quite similar across the different query sets, despite the fact that they represented different dates during the year and different times of day. Perhaps more importantly, the additional information about user click behavior used in the Set 3 results did not result in a substantially different mix of goals. Although this requires further study, it suggests the surprising result that goals can be inferred with almost no information about the user’s behavior.

Because the top level of our goal classification framework is similar to Broder’s web search taxonomy [7], we also examined how the distribution of our queries into the top-level goal categories compared with his. Broder used two methods to classify queries, a user survey and manual classification of log entries. The survey had one question intended to identify navigational queries, and one that allowed users to choose any of several tasks (shopping, downloading, etc.) that he considered “transactional.” If none of these tasks was chosen, the query was assumed to be informational. The log analysis followed a similar decision procedure. Broder also eliminated sexually oriented queries, which accounted for about 10% of the data.

Figure 2 compares our top-level goal classification with results reported by Broder. (We are simplifying somewhat by equating Broder’s “transactional” category with our more general “resource” goal.) We consistently found a greater proportion of informational queries, and a smaller proportion of navigational

¹ The number was not exact because we started with a larger set and then discarded those that were either not English or used non-standard search operators such as “link:”.

Table 3: Results of Classifying Queries by Search Goals

GOAL	SET 1	SET 2	SET 3
directed	2.70%	3.30%	7.30%
undirected	31.30%	26.50%	22.70%
advice	2.00%	2.70%	5.00%
locate	24.30%	25.90%	24.40%
list	2.70%	2.90%	2.10%
informational total	63.00%	61.30%	61.50%
download	4.30%	4.30%	5.60%
entertain	4.00%	8.20%	5.80%
interact	5.70%	4.30%	6.00%
obtain	7.70%	10.30%	7.70%
resource total	21.70%	27.00%	25.00%
navigational	15.30%	11.70%	13.50%

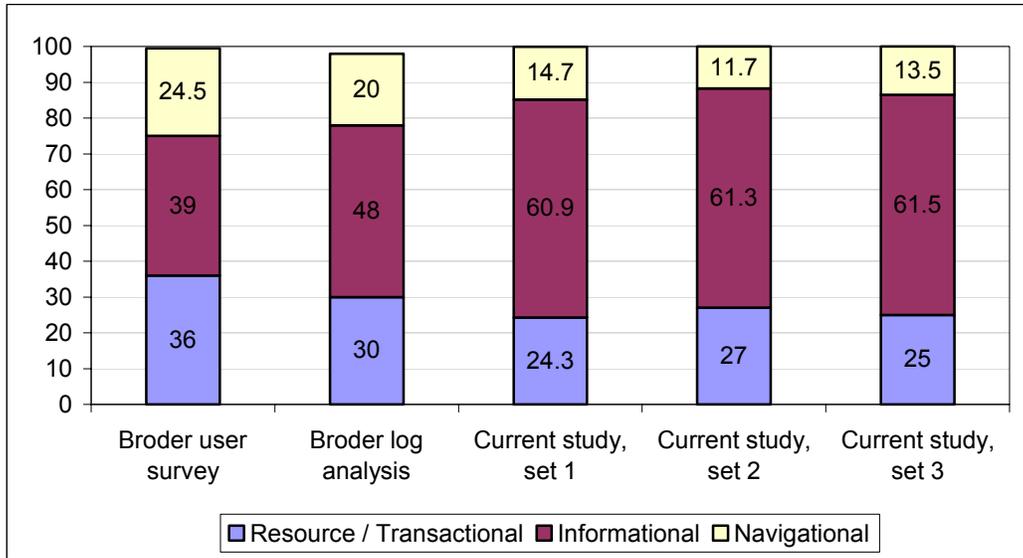


Figure 2: Comparison of Broder's search taxonomy to our top-level goals. Resource and informational results in the first column are Broder's estimates. Results do not total 100% due to rounding error.

and resource/transaction queries than the earlier study. While the differences in informational and resource/transactional queries may be accounted for by our different definitions of those categories, this does not account for the large difference in the fraction of navigational queries.

In fact, since Broder sampled all queries, while we sampled only session-initial queries, the actual difference in navigational query rates may be even higher. This is because navigational query sessions are likely to be shorter and thus overrepresented in our session-initial measure. However, it is not clear that this had any

more impact than the other methodological differences used to obtain our respective data sets.

If our findings about the relatively small number of navigational queries are accurate, they suggest that much of the attention in the commercial search engine world may be misdirected. Tests such as the "Perfect Page Test" organized by one search engine newsletter [12] encourage search engine providers to focus on performance on navigational queries, even though this does not appear to reflect the majority of user needs.

5. FUTURE WORK

In analyzing our results, we are aware of certain limitations that may restrict the generalizability of our conclusions. One issue is that we have no way of knowing conclusively whether the goal we inferred for a query is in fact the user's actual goal. In the future, we would like to combine our work with user studies, including qualitative data such as diary reports of user goals. In order to do this, we first need to make sure that our goal framework and classification methodology can be used by judges other than the authors.

A second issue is that the AltaVista user population may not be representative of search engine users in general. In particular, AltaVista's reputation for providing more powerful query tools, combined with its relatively small market share, may make it the engine of choice for users with difficult informational queries, but not a first choice for typical users issuing common queries. It is possible that this may account for some of the user behavior we saw, despite the fact that we already excluded queries with explicit Boolean syntax or other advanced search operators. In order to investigate this issue, we hope to extend our research to Yahoo! search users.

6. CONCLUSIONS

If web search engines are to continue to improve in the future, they will need to take into account more knowledge of user behavior – not just how people search, but why. We have created a framework for understanding the underlying goals of search, and have demonstrated that the framework can be used to associate goals with queries given limited information.

This analysis of user goals has already yielded two unexpected patterns in web search. First, so-called “navigational” queries appear to be much less prevalent than generally believed. Second, many queries appear to be motivated by a previously unexplored goal involving the need to obtain online and offline resources.

More importantly, an understanding of search goals provides a foundation for tackling the larger problems of conveying user goals to a search engine (either explicitly or by inference), and modifying the engines' algorithms and interfaces to exploit this knowledge.

7. REFERENCES

[1] AltaVista, <http://www.altavista.com>.
[2] AltaVista, description of Prisma query refinement tool, <http://www.altavista.com/help/search/pp>.

[3] Anick, P. Using Terminological Feedback for Web Search Refinement: A Log-Based Study. *Proceedings of SIGIR 2003*, 88-95.
[4] Bates, M.J. Information Search Tactics. *Journal of the American Society for Information Science*, 30, July 1979, 205-214.
[5] Bates, M.J. The Design of Browsing and Berrypicking Techniques for the Online Search Interface. *Online Review* 13, October 1989, 407-424.
[6] Belkin, N.J., Oddy, R.N., and Brooks, H.M. ASK for Information Retrieval: Part II. Results of a Design Study. *Journal of Documentation*, 38(3), Sep. 1982, 145-164.
[7] Broder, A. A Taxonomy of Web Search. *SIGIR Forum* 36(2), 2002.
[8] Google, Description of “I’m Feeling Lucky” feature, <http://www.google.com/help/features.html#lucky>.
[9] Jansen, B.J. and Pooch, U. A Review of Web Searching Studies and a Framework for Future Research. *Journal of the American Society of Information Science and Technology*, 52(3), 235-246, 2000.
[10] Rose, D.E. Reconciling Information-Seeking Behavior with Search User Interfaces for the Web. *Journal of the American Society of Information Science and Technology*, to appear.
[11] Silverstein, C., Henzinger, M., Marais, H., and Moricz, M. Analysis of a Very Large Web Search Engine Query Log. *SIGIR Forum*, 33(3), 1999. Originally published as DEC Systems Research Center Technical Note, 1998.
[12] Sherman, C. and Sullivan, D. The Search Engine ‘Perfect Page’ Test. *Search Day* 391 (Nov. 4, 2002), <http://www.searchenginewatch.com/searchday/02/sd1104-pptest.html>.
[13] Spink, A., Jansen, B.J., Wolfram, D., and Saracevic, T. From E-Sex to E-Commerce: Web Search Changes. *IEEE Computer*, 35(3), 107-109, 2002.