Getting patients to the right healthcare sources:
From real-world questions to *Strategy Hubs*

Suresh K. Bhavnani[1], Christopher K. Bichakjian[2], Jennifer L. Schwartz[2],
Victor J. Strecher[3], Rodney L. Dunn[3], Timothy M. Johnson[2], and Xiaobo Lu[1]

[1]School of Information, [2] Department of Dermatology, [3]Cancer Center
University of Michigan, Ann Arbor MI, 48109-1092

ABSTRACT

To address the growing problem of novices accessing incomplete and unreliable information on the Web, this paper begins by probing the nature of strategic knowledge known by expert healthcare searchers. The analysis reveals the existence of a strategy that is focused to a specific type of healthcare question. To investigate whether there exist other similar strategies, we conduct two pilot studies. The first leads to the development of a taxonomy for skin cancer questions based on real-world questions. This taxonomy has high inter-rater agreement when tested with new real-world questions. The second pilot helps to identify 6 strategies that are specific to question types in the taxonomy. We demonstrate how these strategies can be made available to a wide range of users through a prototypical system referred to as a Strategy Hub. In conclusion, we briefly describe our current and future research to test if such a system can make patients more effective and efficient in the retrieval of reliable and relevant healthcare information on the Web.

INTRODUCTION

Despite huge advances in making healthcare information accessible to vast numbers of users, the effective retrieval of relevant and reliable healthcare information remains a challenge. While healthcare information seeking has become the most common task performed by Web users [1], there has been a corresponding increase in the number of scientifically unreviewed websites that provide out-of-date, and incorrect healthcare information [e.g. 2, 3]. Because most healthcare information seekers rely on general-purpose search engines [1] like Google that do not distinguish between reliable and unreliable sources, many searches are ineffective, and potentially dangerous.

Although most users are novices in searching for healthcare information, professionals such as medical reference librarians and informaticians have learned how to search the Web using sophisticated strategies. However, few studies have analyzed the knowledge that such healthcare search experts use while actually performing search tasks. What is the nature of the strategic knowledge that such experts have, and how can such knowledge be made available to patients?

To address this question, we began by probing the nature of knowledge used by a medical reference librarian with many years of experience in searching for healthcare information. A goal decomposition of her behavior revealed a search strategy that was specific to the type of

question she was answering. This strategy was qualitatively different from the goal sequence of a novice who used a general-purpose search engine to answer the same question.

The existence of a search strategy that was specific to a question type alerted us to the importance of such knowledge because general-purpose search engines like Google are not designed to provide such knowledge. We therefore performed two pilot studies to explore if there were other similar question-specific strategies. The first pilot helped to develop a taxonomy of real-world skin cancer questions, and the second helped to identify strategies tailored for specific question types in the taxonomy. These pilots motivated us to develop a prototype for a new kind of website called a *Strategy Hub* which demonstrates how such strategies could be made available to novice users searching for healthcare information. We discuss the insights we gained from these pilots, and conclude by briefly describing our future research to empirically test the Strategy Hub with real patients searching for online healthcare information.

ANALYSIS OF HEALTHCARE SEARCHING BEHAVIOR

In a recent study [4], we compared the search behavior of expert medical reference librarians, to users who had equivalent experience in using Web browsers and search engines, but negligible experience in searching for healthcare information. The study showed general patterns of behavior across the participants. Our aim in this section is to provide a detailed understanding into the specific nature of the knowledge exhibited by an expert and a novice healthcare searcher from the above study. Each attempted to retrieve information from the Web for the following healthcare question:

Tell me three categories of people who should, or should not get a flu shot and why?

To answer the above question, the reference librarian first accessed MEDLINEplus[1], a reliable healthcare collection. Next, she used the query "flu shot" to search within MEDLINEplus and found several sources of information within that collection. She visited two of those sources and retrieved categories of people who should and should not get a flu shot. Not being satisfied with the sources she had visited, she retrieved the name of a flu shot ("Flushield") from a third link because she had the explicit goal of

---

[1] http://www.medlineplus.gov

verifying the information she had obtained from MEDLINEplus.

She then attempted to verify the information obtained through MEDLINEplus by visiting a pharmaceutical company that sold Flushield. She did this by first visiting the site rxlist.com but failed to find the information. She then used Google[2] to find a pharmaceutical company that sold Flushield (wyeth.com), and verified the information she had obtained from MEDLINEplus by reading the indications and contraindications for the vaccine.

Although she did not explicitly search for more categories than required by the task, she completed the task by having access to a comprehensive list of 9 categories of people who should get a flu shot, and 5 categories of people who should not. She took approximately 7 minutes to complete the task, and visited 3 reliable healthcare sources.

Figure 1 shows a hierarchical goal decomposition [5] of the above behavior. This analysis method progressively decomposes a given goal (in this case the question) starting from the question at the top of the decomposition, to any level of detail required by the analysis. As shown, the decomposition reveals a critical strategy with the steps: (1) find flu shot information source in reliable collection, (2) find categories within reliable source, and (3) verify categories in pharmaceutical website. These three steps occur early in the goal decomposition and are specific to the drug-related question being addressed.

To use this strategy, the user must know two types of knowledge. (1) *Declarative knowledge* (facts and concepts) of the distinction between reliable healthcare collections such as those sponsored by governments and universities, and unreliable sites such as personal pages, in addition to their URLs. (2) *Procedural knowledge* of how to select and sequence the declarative knowledge for specific tasks.

When the same task was performed by the novice healthcare searcher (who had equivalent experience in searching for information on the Web, but negligible experience in searching for healthcare information), he did not exhibit the strategy used by the expert. Instead, he went directly to Google and typed "who should or should not recieve (sic) flu shots" in the query box. He visited numerous hits provided by Google *in roughly the same order* of the displayed hits. He took approximately 20 minutes to complete the task, and visited 13 sites suggested by Google, none of them high-quality healthcare sites recommended by the Consumer and Patient Health Information Section of the Medical Library Association.

The novice's main strategy was to use a general-purpose search engine. Such engines are neither designed to provide the knowledge of which sites are reliable, nor are they designed to provide the procedural knowledge of how to sequence the different stages of the strategy. Users who rely on Google therefore can obtain knowledge about

Tell me three categories of people who should,
or should not get a flu shot and why?
   Find flu shot information source in reliable collection
      Search MEDLINEplus.gov
         Go to MEDLINEplus.gov
         Search within MEDLINEplus.gov
   Find categories within reliable source
      Search page in MEDLINEplus.gov
         Search source from MEDLINEplus.gov
         Search page
   Verify categories in pharmaceutical website
      Search rxlist.com
         Go to rxlist.com
         Search rxlist.com
      Search Google.com
         Go to Google.com
         Search within Google.com
      Search wyeth.com
         Go to wyeth.com
         Search page

Figure 1. Goal decomposition of the flu-shot task as performed by the expert. Bolded text shows the question-specific strategy used by the expert to complete the task.

URLs, but not the selection and sequencing knowledge known by the expert. Because such knowledge cannot be easily inferred from the information provided by Google, the novice retrieved fragmented pieces of information from a variety of unreliable sites in almost thrice the amount of time. The behavior of the expert and novice described above were typical of other participants in the study [4].

Given the importance of such question-specific strategies to the search behavior described, we were motivated to explore if there were other question-specific strategies in healthcare. To identify such strategies, we first needed to develop a taxonomy of real-world questions reflecting the healthcare information needs of real users.

PILOT-1: DEVELOPMENT OF A TAXONOMY FOR SKIN CANCER QUESTIONS

Our research team included dermatologists specializing in skin cancer who had recently conducted and published an extensive study on the quality of melanoma information on the Web [3]. We therefore focused on developing a taxonomy of skin cancer questions.

Method for building a question taxonomy: To understand the needs of real users, we accessed the database of a well-known *Ask-a-Doc* website called NetWellness[3] [6]. Ask-a-Doc sites enable users to send anonymous healthcare questions in full-text, which are then replied to by a physician who specializes in the area of the question. The question-answer pairs are stored in a database, and are publicly accessible through a search engine. NetWellness is one of the oldest Ask-a-Doc sites and currently contains over 17,000 healthcare question-answer pairs on a wide range of diseases.

On January 13, 2002, we attempted to retrieve all skin cancer questions from this database by entering the following terms: melanoma, skin cancer, basal cell carcinoma, basal cell cancer, squamous cell carcinoma,

---

[2] http://www.google.com

[3] http://www.netwellness.org

squamous cell cancer, BCC, SCC, and nonmelanoma skin cancer. These terms had to be entered separately because the search engine did not handle Boolean expressions consistently. The resulting overlapping hits were removed. The above method retrieved 216 question-answer pairs.

Because the search engine also did not distinguish whether the query terms occurred in the question, or in the answer, we asked a cancer expert to remove from this set all questions not relevant to skin cancer. Furthermore, duplicate questions with identical wordings, and follow-up questions that referred to earlier questions were also removed. The resulting set consisted of 61 questions that directly related to skin cancer.

The above set was randomly divided into two roughly equal sets: Set A consisted of 30 randomly selected questions that were used by two skin cancer experts (who had not seen the questions before) to develop a question taxonomy. Set B consisted of 31 randomly selected questions that were used in a subsequent inter-rater reliability test.

Each of the questions in Set A was printed on a card, and the skin cancer experts were asked to categorize them according to how they saw fit. If a question had more than one part, then only the first part was used for the categorization. The experts spent 3 sessions of about 2 hours each to develop the taxonomy. The researcher encouraged both the experts to discuss their rationale for the question categories, and formalized their categorization schemes between each session. For example, when the experts were categorizing the questions in terms of a hierarchy, the researcher formalized the hierarchy on large poster boards for the subsequent session in order to assist in the categorization.

Results: The above process resulted in a hierarchical taxonomy that started at the level of skin-related conditions, and became progressively more specific. The first column of Table 1 shows the top levels of the taxonomy[4].

Because we wished to focus only on skin cancer questions, that subset was further categorized as shown in the first column of Table 2. The first level of the skin cancer hierarchy consisted of disease categories. This consisted of Skin Cancer, Melanoma, Basal Cell Carcinoma (BCC), Squamous Cell Carcinoma (SCC), and Other (which contained questions related to rare types of skin cancer like Merkel cell carcinoma). The questions at this level consisted of general requests for information about a skin cancer. The next level of the taxonomy consisted of the 5 well-known medical categories: Terminology, Risk/Prevention, Diagnosis, Treatment, and Prognosis. For

---

[4] Despite having the cancer expert remove all non-skin cancer questions from the original set of 216 questions, the skin cancer experts in the experiment were more discriminating of what was *not* a skin cancer question. Table 1 shows how the skin cancer experts separated the non-skin cancer questions from the skin cancer questions. Throughout the pilot, the experts chose to create categories even when there were no questions to put in them, in order to create a more complete taxonomy.

Question Taxonomy	Definition	%
Skin	General question related to skin	2
Non Disease	Related to skin but not related to a skin disease	10
Disease	Related to a specific skin disease	8
Pre-Skin Cancer	Related to any pre-cancerous skin disease	0
Melanoma	Pre-cancerous lesions related to melanoma	3
BCC	Possible, but incorrect question related to BCC	0
SCC	Pre-cancerous lesions related to SCC	3
Skin Cancer	All skin cancer questions shown in Table 2	57
Non Skin	Not related to skin or skin disease	15
Unable to categorize	Does not fit in any category	2

Table 1. Taxonomy for skin cancer questions obtained from NetWellness. The taxonomy for Skin Cancer questions (marked in gray) is exploded in Table 2. Percentages are based on all questions in Set A and B (rounded to the nearest percent).

example, the following question was classified as a Treatment question for BCC: *I have...a small Basal Cell nonmelanoma cancer on the side of my face...I have been advised to have Mohs Micrographic Surgery...do you feel that Mohs Micrographic Surgery is my only treatment?* The next level of the taxonomy dealt with specific issues related to each of the above categories, as described in the table. For example, the following question was classified as a Risk/Prevention-Statistical question for melanoma: *Do African Americans suffer from melanoma or other forms of skin cancer caused by sun damage?*

Method for test of inter-rater reliability: An inter-rater

Question Taxonomy	Definition	%
Skin cancer type	General question related to skin cancer type	3
Terminology	Broad terminology question	0
Comparative	Comparison of two skin cancer terms	0
Definition	Definition of a skin cancer term	9
Association	Relationship between two or more skin cancer terms	6
Risk/Prevention	General question about risk or prevention	0
Statistical	Statistics about risk or prevention	9
Qualitative	Qualitative issues related to risk or prevention	11
Specific	Role of a specific risk or prevention factor	11
Diagnosis	General question related to skin cancer diagnosis	0
Self-examination	Detection of skin cancer based on a visual description	20
Doctor's-exam	Issues related to a doctor's examination of skin cancer	0
Test	Issues related to the procedure or results of a skin cancer test	9
Treatment	General question related to skin cancer treatment	11
Stages	Broad issues related to treatment stages	0
Stage-1	Information related to stage-1 treatment	0
Stage-2	Information related to stage-2 treatment	0
Stage-3	Information related to stage-3 treatment	0
Stage-4	Information related to stage-4 treatment	6
Clinical-trial	Information related to clinical trials	6
Prognosis	General question related to skin cancer prognosis	0
Statistical	Statistics about prognosis	0
Qualitative	Qualitative issues related to prognosis	0

Table 2. Taxonomy for skin cancer questions obtained from NetWellness. The above taxonomy repeats for melanoma, BCC, SCC, and other rare types of skin cancers such as Merkel cell carcinoma. Percentages are based on total skin cancer questions in Set A and B (rounded to nearest percent). The categories in gray represent 60% of the total skin cancer questions.

experiment was conducted to test if two skin cancer experts could reliably categorize the questions in Set B. The 31 questions from Set B were printed on separate 2"x2" cards, and the hierarchical taxonomy, shown in Table 1 and 2 was drawn on 2'x3' poster boards. The two experts were asked to independently categorize the questions using the taxonomy, and to speak aloud as they performed the task. The sessions were video taped, and the categorizations recorded.

Results: The two experts agreed completely on 25 of the 31 questions (80.6%). Of the 6 questions they differed on, 3 were different by one level of categorization (2 Treatment vs. Treatment-Stages, 1 Skin Cancer-Risk/Prevention-Statistical vs. Skin Cancer-Melanoma-Risk/Prevention-Statistical) and 3 differed by non-skin cancer category (2 Skin-Disease vs. Skin-Non Disease, 1 Non Skin vs. Unable to Categorize). To test the level of agreement, the sub-nodes of the skin cancer taxonomy were collapsed to the second level (Terminology, Risk/Prevention, etc.) so that different levels would be independent. After collapsing, the experts agreed on the categorization of 27 of the 31 questions (87.1%), which had a level of agreement of 0.84 (Cohen's Kappa, 95% CI 0.70-0.99).

The pilot therefore enabled us to develop a question taxonomy for skin cancer that was based on real-world questions. Furthermore, this taxonomy had high inter-rater agreement for the categorization of new questions.

PILOT-2: IDENTIFICATION OF STRATEGIES THAT ARE SPECIFIC TO QUESTION TYPES

In our second pilot study, our goal was to explore if there existed a set of strategies that were specific to the question types in the taxonomy. For the pilot, we chose to focus on the most frequent question categories. As shown in Table 2, the categories for Risk/Prevention and Diagnosis formed 60% of the total skin-cancer questions.

The questions in each of the above categories were generalized to the form: *My question relates to <disease topic> for <disease type>*. We refer to this generalized form as a *question type*. For example, the melanoma questions in the Self-examination category were generalized to the question type: *My question relates to <self-examination in the diagnosis> for <melanoma>*. The goal of this pilot was to identify strategies that were specific to 6 question types for the sub-nodes of the most frequent categories: Risk/Prevention and Diagnosis. We also limited the strategies to melanoma as we had experts who specialized in that specific form of skin cancer.

Method: Two skin cancer experts were given the 6 question types described above, and were asked to pool their past experience to describe explicitly the steps they would take to answer each question. They were encouraged to access the Web to identify sites appropriate to retrieve information for each step.

Results: Table 3 shows four of the six strategies that were identified by the skin cancer experts. Each strategy

Question types	Strategies specific to question type
1. My question relates to *statistical information on risk factors and prevention for melanoma*	Step 1: Understand general cancer statistics Step 2: Obtain general melanoma statistics Step 3: Obtain detailed melanoma statistics
2. My question relates to *qualitative information on risk factors and prevention for melanoma*	Step 1: Learn about melanoma prevention Step 2: Learn about melanoma risk factors Step 3: Estimate your risk of melanoma
3. My question relates to a *specific risk factors on risk factors and prevention* for *melanoma*	Step 1: What are the effects of ultraviolet radiation? Step 2: Learn about how to reduce exposure to ultraviolet radiation Step 3: Learn about risks of tanning booths
4. My question relates to *self-examination in the diagnosis* for *melanoma*	Step 1: Learn about the ABCDs of moles and melanoma Step 2: Understand the difference between moles and melanoma Step 3: Learn how to do a self skin examination Step 4: Locate a dermatologist near you

Table 3. Four of the six strategies identified by the skin cancer experts to structure the search of information for specific question types in the taxonomy. Each step in a strategy is associated with one or more URLs (not shown) to achieve the step.

consists of steps with associated URLs (not shown in the table) to find that information on the Web. The strategies appear to have three important characteristics. (1) Each of them is specific to a question type. However, some strategies appear to be easily modified to be useful for other diseases. For example, the strategy related to finding statistical information for melanoma (question-type 1) could be modified to be useful for any cancer by replacing the URLs for each step. (2) Some steps in the strategies appear to be based on medical subject knowledge, while others appear to be based on knowledge of new types of information on the Web. For example, Step 3 (Learn how to do a self skin examination) in question-type 4 is typical of medical advice given by a physician to a patient. In contrast, Step 3 (Estimate your risk of melanoma) for question-type 2 is predicated on the existence of a new set of tools called *risk calculators*. In this case, there exists a calculator in the Harvard medical site[5] that estimates a user's risk for developing melanoma based on a series of questions. Knowing of the existence of such an online tool is the main rationale for the step in this strategy. (3) The URLs associated with each step included government and educational sites, in addition to commercial sites. This demonstrated that experts do not outright reject commercial sites, but know which commercial sites are reliable. This is counter to the general claim that only government (.gov), education (.edu), and organization (.org) sites are reliable sources for healthcare information.

Pilot-2 therefore demonstrated the existence of strategies that are specific to question types, some of which can be modified to be generally useful across certain diseases. Future research will enable us to understand the range of generality for such strategies.

---

[5] http://www.yourcancerrisk.harvard.edu/hccpquiz.pl?func=start&quiz =melanoma

DESIGN IMPLICATIONS: *STRATEGY HUBS*

The above pilot demonstrates the existence of search strategies that are specific to question types. This is also precisely the knowledge that is neither available from the general-purpose search engines typically used by novice healthcare searchers, nor through domain-specific portals such as MEDLINEplus. We were therefore motivated to develop a system to deliver such strategic knowledge to novice users. Towards that goal, we have developed a prototype for a new class of websites called a *Strategy Hub*. The goal of a Strategy Hub is to provide the declarative and procedural knowledge used by experts to search for information in a specific domain such as healthcare.

The Strategy Hub prototype allows a user to select a disease, and a disease topic. For example, a user can select *Melanoma*, followed by selecting *Statistical* under the node *Risk/Prevention* (representing question-type 1 in Table 3). The system responds by displaying the corresponding steps of a strategy as shown in Figure 2. When the user selects one of the URLs in a step, the lower frame displays the contents of that webpage. The user can now find information in that page before proceeding to the next step in the strategy. All the steps of the strategy are visible throughout the session enabling the user to always be aware of the context of the entire strategy. Our future research will refine the above prototype and test it with skin cancer patients from the University of Michigan Department of Dermatology.

CONCLUSION AND FUTURE RESEARCH

We have described our initial steps in understanding the nature of knowledge used by healthcare search experts, and an approach to make that knowledge available to large numbers of users. A goal decomposition of the search behavior of a healthcare search expert revealed a strategy that was specific to a question type. This knowledge is difficult to infer from the information provided by general-purpose search engines which novice healthcare searchers typically use. To understand if there were other similar strategies that were specific to question types, we conducted two pilot studies. The first identified a taxonomy of real-world questions related to skin cancer. The second pilot study identified 6 strategies that were specific to question types in the skin cancer taxonomy. The pilot studies provided the motivation to develop a prototype for a new kind of website called a Strategy Hub which provides users with information search strategies that are tailored to specific question types.

Our future research includes the development of similar taxonomies of real-world questions for different types of cancers, and to identify strategies that are specific to question types in the taxonomy. We are acutely aware of the time-consuming nature of our approach for identifying strategies to service the large number of diseases in healthcare, and we are exploring more efficient means to operationalize this knowledge. However, we believe many of the strategies will fall into general patterns that should
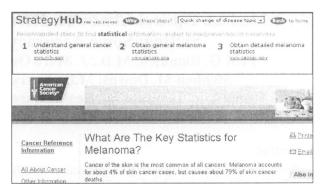

Figure 2. Interface of the Strategy Hub for healthcare showing the steps of a strategy specific to a question type. Each step has an associated URL, which, when selected, is displayed in the lower frame. The user therefore has access to the declarative and procedural knowledge known by expert healthcare searchers.

enable them to be reused across diseases with minor modifications. We also plan to test the hypothesis that the declarative and procedural knowledge provided by Strategy Hubs will enable novice healthcare searchers to be more effective and efficient in searching for healthcare information compared to the current approach of using general-purpose search engines.

ACKNOWLEDGEMENTS

REFERENCES

1. Fox, S., Rainie, L. The online health care revolution: How the Web helps Americans take better care of themselves. Pew Internet and American live project: Online life report. Available from: URL: http://www.pewinternet.org/reports/toc.asp?Report=26
2. Biermann, S., Golladay, G., Greenfield, M., Baker, L. Evaluation of cancer information on the Internet. Cancer, 1999;86(3):381-91.
3. Bichakjian, C., Schwartz, J., Wang, T., Hall J., Johnson, T., Biermann, S. Melanoma information on the Internet: Often incomplete-a public health opportunity? Journal of Clinical Oncology, 2002;20(1):134-41.
4. Bhavnani, S. Important cognitive components of domain-specific search knowledge. In: Voorhees, EM, and Harman, DK, editors. Proceedings of the Text Retrieval Conference (TREC); 2001 Nov 13-16; Gaithersburg, Maryland. Gaithersburg: Maryland; 2001. p. 571-8.
5. Kirwan, B., Ainsworth, L. K., editors. A guide to task analysis. London: Taylor & Francis, 1993.
6. Morris, T., Guard, R., Marine S., Schick, L., Haag, D., et al. Approaching equity in consumer health information delivery: NetWellness. Journal of the American Medical Informatics Association. 1997;4(1):6-13.