Strategy Hubs: Next-Generation Domain Portals with Search Procedures

Suresh K. Bhavnani[1], Christopher K. Bichakjian[2], Timothy M. Johnson[3], Roderick J. Little[4], Frederick A. Peck[5], Jennifer L. Schwartz[6], Victor J. Strecher[7]

School of Information[1, 5] University of Michigan Ann Arbor, MI 48109-1092

Department of Dermatology[2, 3, 6] University of Michigan Ann Arbor, MI 48109-1092

Cancer Center[4, 7] University of Michigan Ann Arbor, MI 48109-1092

{bhavnani[1], chriskb[2], timjohn[3], rlittle[4], peckf[5], jennschw[6], strecher[7]}@umich.edu

ABSTRACT

Current search tools on the Web, such as general-purpose search engines (e.g. Google) and domain-specific portals (e.g. MEDLINEplus), do not provide search procedures that guide users to form appropriately ordered sub-goals. The lack of such procedural knowledge often leads users searching in unfamiliar domains to retrieve incomplete information. In critical domains such as in healthcare, such ineffective searches can have dangerous consequences. To address this situation, we developed a new type of domain portal called a *Strategy Hub*. Strategy Hubs provide the critical search procedures and associated high-quality links that enable users to find comprehensive and accurate information. This paper describes how we collaborated with skin cancer physicians to systematically identify generalizeable search procedures to find comprehensive information about melanoma, and how these search procedures were made available through the Strategy Hub for healthcare. A pilot study suggests that this approach can improve the efficacy, efficiency, and satisfaction of even expert searchers. We conclude with insights on how to refine the design of the Strategy Hub, and how it can be used to provide search procedures across domains.

Keywords: Strategy Hub, healthcare, Web searching.

INTRODUCTION

Millions of searchers regularly use modern search engines like Google to find information on the Web. The goal of such search engines (in the words of the developers) is to "get you to the right site" [15]. These search engines are effective to find information for questions such as "What is a melanoma?" that have *specific* answers. However, current search engines are far less effective when users need to find information for questions that require a *comprehensive* understanding of a topic such as "What are the treatment options for Stage III melanoma?" This is because no one source contains all the information relevant to get a comprehensive understanding of treatment for Stage III melanoma. For example, while the National Cancer Institute site provides primary treatment information about melanoma, supplemental treatments such as the use of *interferon* specifically for melanoma, are described in other sources like the University of Michigan's Cancer site. To get a comprehensive understanding of treatment for Stage III melanoma, users must therefore first have the sub-goal to retrieve *primary* treatment information about the disease, followed by a second sub-goal to look for *supplemental* treatments in a specialized source.

CHI 2003, April 5–10, 2003, Ft. Lauderdale, Florida, USA.

Such ordered sub-goals are neither spontaneously obvious from the hits provided by Google, nor from the coarse-grained taxonomies in domain portals such as MEDLINEplus, a leading healthcare portal used by search experts [1, 2]. For example, the melanoma page in MEDLINEplus, provides three links under the heading "Treatment", none of which point to supplemental treatments. As a result, users often retrieve incomplete information [2, 4]. In critical domains such as healthcare, such searches can have dangerous consequences.

To address the above inadequacies of current search tools, this paper describes the development and testing of a new type of domain portal called a *Strategy Hub*. Strategy Hubs can be regarded as the next step in the evolution of domain portals and are designed to provide search procedures that make explicit the critical sub-goals in a particular order to enable users to get a comprehensive understanding of a topic.

We begin by describing the factors that make search procedures critical to get a comprehensive understanding of a topic. We then describe how we used an existing empirically-based taxonomy of skin cancer questions to identify expert search procedures to find comprehensive information about melanoma, a deadly form of skin cancer. Next, we describe the design of the Strategy Hub to provide these search procedures, and a pilot experiment to compare its performance with conventional search tools. We conclude with ideas of how to improve the design of the Strategy Hub, and how it can be used to provide search procedures across domains.

THE IMPORTANCE OF PROCEDURAL SEARCH KNOWLEDGE

Previous research has noted the critical role of procedural search knowledge when searching for comprehensive information about a topic [1, 2, 11]. For example, a recent

study [2] described the three-step search procedure followed by an expert healthcare searcher looking for flu shot information. (1) Access a reliable healthcare portal such as MEDLINEplus to identify sources for flu-shot. (2) Access a high-quality source of information to retrieve general flu shot information. (3) Verify the information by visiting a pharmaceutical company that sells a flu vaccine. Such search procedures enabled experts to find comprehensive information quickly and effectively compared to novices who were unable to infer such knowledge by just using Google [1].

The above study suggests that users who have acquired search expertise in a particular domain know more than just how to use query-based search engines, and the names of high-quality sources [13]; they have also acquired the *procedural search knowledge* to determine which sites to visit in which order when searching for comprehensive information. This procedural search knowledge consists of three components. (1) The sub-goals to organize a search in a particular domain. For example, the healthcare search expert knew the critical sub-goal of *verifying* healthcare information by visiting a pharmaceutical source. (2) The order in which to satisfy those sub-goals. For example, the expert knew to first visit a general domain portal to get broad and general information, before visiting a pharmaceutical company to verify that information. (3) The selection knowledge to decide which sites or pages will satisfy a specific sub-goal, such as to visit MEDLINEplus to obtain reliable healthcare sources.

But why has the above situation emerged? What are the factors that make procedural search knowledge to visit sites in a particular order so important?

Variability in detail, accuracy, and completeness

The rapid proliferation of websites has resulted in an unprecedented variability in the nature of information available. No one site provides all levels of detail, accuracy, and completeness of information about most complex topics. For example, some healthcare sites provide broad overviews of healthcare topics useful for introductions, while other sites provide more specific details about new treatments and still others provide inaccurate, incomplete, and contradictory information [e.g., 4] for some topics, but not for other topics.

Given this variability, experts develop sub-goals to organize a search based on their understanding of the information sources available. As described in the introduction, a comprehensive understanding of Stage III melanoma requires users to first visit a site providing general information, followed by visiting a different site that provides details about supplemental treatment. Because this procedural search knowledge is not explicitly provided by current search tools, novices must infer important sub-goals through trial and error.

Specialization of content in websites

Another reason why procedural knowledge has become critical when searching for comprehensive information about a topic is how sites within a domain have specialized into various site genres [6]. For example, in the healthcare domain, sites range from *ask-a-doc* sites (that provide answers to healthcare questions from a real doctor), to sites that provide *risk calculators* (that calculate your risk for suffering from a particular disease). E-commerce sites range from *review* sites (that provide reviews by consumers for different products), to *price comparison* sites (that provide a list of online vendors that sell a product, ranked by price).

Experts in various domains know about the existence of these site genres and form sub-goals to exploit this specialization when searching for information. For example, in the study mentioned earlier [2], shopping experts, when looking for three low prices for a new digital camera, first visited review sites (e.g. cnet.com) to learn which cameras were highly rated, followed by finding low prices through a price comparison site (e.g. mysimon.com), followed by looking for discounts in yet another set of sites that advertised online discounts (e.g. staples.com). Once again the experts had procedural knowledge of how to sequence their search through genres of sites leading to a *comprehensive* understanding of digital cameras and their prices. Such searches led to lower prices for high quality cameras when compared to novices who relied on Google for their searches, and who did not infer the critical sub-goals known by the experts from the links provided by Google.

The need to make explicit procedural search knowledge

Although experts have identifiable search procedures that appear to improve search efficiency and effectiveness, such knowledge is not easily inferred from the information provided by conventional search tools. General-purpose search tools like Google provide a ranked list of URLs that are relevant to the query based on link analysis [16], and the occurrence of the query in the pages. While the ranking algorithm attempts to give higher ranks to pages that are most pointed to, there is no explicit guarantee for the reliability of such highly ranked sites, nor any indication of the critical sub-goals in a domain to guide which pages to visit in which order. As described earlier, domain portals such as MEDLINEplus do provide reliable sources of information, but also do not provide the procedural knowledge to organize a search. This focus on declarative knowledge, but with the absence of procedural knowledge, directly affects the performance of users when they search in unfamiliar domains.

Because such search procedures are neither provided by general purpose search engines, nor by domain portals such as MEDLINEplus, novices must infer the important sub-goals, order them, and determine which URLs to visit to satisfy each sub-goal. This process can be error-prone and time-consuming. We therefore were motivated to see if we could systematically identify critical search procedures, and make them available to novice users.

IDENTIFICATION OF SEARCH PROCEDURES TO FIND COMPREHENSIVE INFORMATION ABOUT MELANOMA

We chose to focus our initial research on melanoma (a deadly form of skin cancer) because we had access to skin cancer physicians who specialized in treating melanoma. Furthermore, besides having medical knowledge, the physicians had recently completed a survey of melanoma information on the Web [4], and therefore also had search expertise. Although we choose to focus our research on the identification of search procedures for melanoma, we will show how they generalize to other sub-domains in healthcare and beyond.

To begin a systematic identification of search procedures, we developed an empirically-based taxonomy of real-world skin-cancer questions [3]. This taxonomy, developed by skin cancer physicians, was based on real-world questions, and had high inter-rater reliability. The top-level nodes in the taxonomy are similar to those identified by Pratt et al. [12], with the exception of the *terminology* node identified by the physicians in our study. The first column in Figure 1 shows this skin cancer taxonomy with three of the five high-level nodes decomposed to their respective leaf nodes.

To assist the physicians in identifying search procedures for leaf nodes in the taxonomy, we generalized each leaf node in the taxonomy to the form, *My question relates to <topic in the taxonomy> for <disease in the taxonomy>*. This generalized form is referred to as a *question type*. For example, the melanoma questions in the Risk/Prevention Qualitative category were generalized to the question type: *My question relates to <qualitative information on risk factors and prevention> for <melanoma>*. Examples of these question types are shown in the second column of Figure 1.

Two skin cancer experts were given 15 question types based on the 15 leaf-nodes in the taxonomy, and were asked to pool their past experience to describe explicitly the steps they would take to answer each question type. Furthermore, they were encouraged to access the Web to identify webpages that were appropriate to retrieve information for each step. This resulted in the identification of 15 search procedures, one for each node in the taxonomy (three of which are shown in the third column of Figure 1). Each search procedure consists of two to four steps, which represent critical sub-goals to find information about a topic. Analysis of the 15 search procedures to find comprehensive melanoma information helped to identify two important characteristics:

(1) The search procedures exploit the variability in the detail and specialization of webpages. For example, step 3 (estimate your risk for melanoma) for search procedure A in Figure 1 exploits the existence of a melanoma risk calculator on Harvard's cancer-prevention website. However, this site does not provide a description of the risk factors that are used in the estimate and therefore users must first visit other webpages that provide this information. This is reflected in steps 1 and 2 of the same procedure, each of which requires the user to visit different links.

(2) Although the procedures identified by the skin cancer physicians were focused on searching for melanoma information, they followed patterns that generalized into six search procedure *templates*. The fourth column in Figure 1 shows two such templates that appear useful for other diseases within healthcare, and for other domains. For example, the *problem-solution* template could be used to identify search procedures for other diseases such as HIV/AIDS (1. Understand the nature of the HIV retrovirus, 2. Learn about antiretroviral drug therapy, 3. Learn about

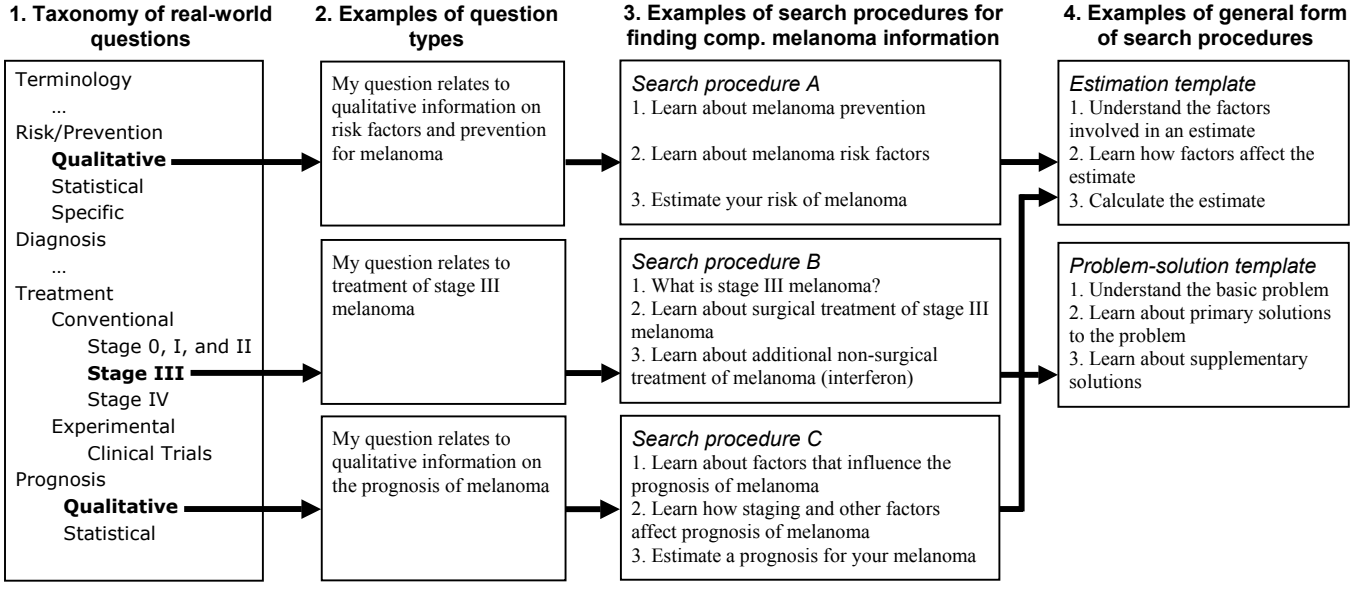

Figure 1. A taxonomy of real-world questions (Column 1), was abstracted to question types (Column 2). The question types were used by skin cancer physicians to systematically identify search procedures (Column 3), which were subsequently generalized into templates (Column 4). These generalized templates are useful to identify search procedures in other healthcare sub-domains like HIV/AIDS, and other online domains like software engineering.

immune-based therapies), and to other domains such as software engineering (1. Learn about database integration with ColdFusion, 2. Learn general methods of displaying database content on a ColdFusion page, 3. Find custom tags to display the data in specific formats). We are currently using these generalized templates to rapidly identify search procedures for other diseases in the healthcare domain.

Having identified expert search procedures for melanoma, we were motivated to explore how such search procedures could be made available on the Web in a new form of domain portal called a *Strategy Hub*.

DESIGN AND DEVELOPMENT OF THE STRATEGY HUB

As described earlier, neither search engines, nor domain portals provide the search procedures that we have identified from search experts. Therefore, users searching for comprehensive information in an unfamiliar domain have the difficult task to infer these search procedures from a list of ranked hits, or from coarse-grained selection categories typically provided by domain portals. We therefore designed a new kind of domain portal called a Strategy Hub to address this issue.

Strategy Hubs have two characteristics that distinguish them from conventional portals. (1) They provide selection categories that are defined at a finer-grained level to enable users to learn more precisely how information in the domain is organized, and to select appropriate topics of interest. For example, while "Treatment" is a leaf node in the categories provided by MEDLINEplus, our implementation of a Strategy Hub provides two more levels of specificity below Treatment as shown in Figure 2. (2) They provide explicit search procedures consisting of ordered sub-goals, in addition to reliable links to satisfy each sub-goal to find comprehensive information about a selected topic. For example, selection of the node Stage III in Figure 2 will provide the search procedure shown in Figure 3.

The design of the Strategy Hub prototype was guided by design principles critical for search interfaces that have been suggested by Shneiderman et al. [14], Furnas et al. [10], and Egan et al. [8]. We first focus on three principles that were critical for the success of the Strategy Hub, and then briefly discuss how the other principles guided the design of our prototype.

(1) Reduce the cost of mapping a question to a node in the taxonomy. As shown in Figure 2, the home page guides the user to select a disease, and a disease topic from a hierarchy based on the empirically based taxonomy developed by the experts. The hierarchies behave similarly to a directory structure in Windows where nodes can be opened up to the leaf nodes. While such hierarchies are intuitively clear, several studies have shown the difficulty that users have in mapping real-world goals to interface elements such as icons on an interface. Furnas et al. [10] refer to this as the *vocabulary problem*, which is based on the observation that users differ substantially in the terms they use to describe a goal or object, and the overlap of the terms between users is small. The vocabulary problem therefore leads to the difficulty of providing short descriptions of interface elements, such as in a taxonomy, which serves all users.

Furnas et al. suggest that the vocabulary problem can be addressed by providing multiple synonyms for interface elements. As shown in Figure 2, this problem is addressed by adding the *Brief explanation* box in the interface of the Strategy Hub, which provides explanations in grammatical sentences that increase the overlaps between the terms people use to describe concepts. In addition to providing explanations, the brief explanation box also provides the steps of the procedure to give the user some *information scent* [5] about the steps to expect. The brief explanation box reduces the chance of users selecting the wrong nodes.

(2) Provide focus plus context. Once the user selects a disease and disease topic, the system responds by providing a sequence of recommended steps to search for information related to the disease topic, in addition to links at each step. As shown in Figure 3, the steps and links are displayed in the upper frame of a dual frame design. When links in the upper frame are selected, the associated page is displayed

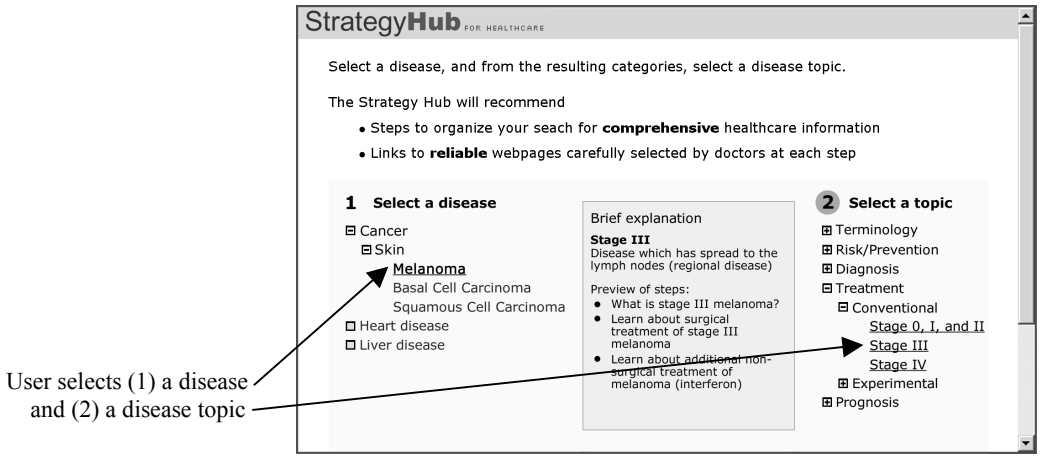

Figure 2. The home page of the Strategy Hub that guides the selection of a disease (melanoma), and a disease topic (Stage III). The *Brief explanation* box provides explanations when the cursor is placed above a particular node, in addition to a preview of the steps of a search procedure to find information related to that node.

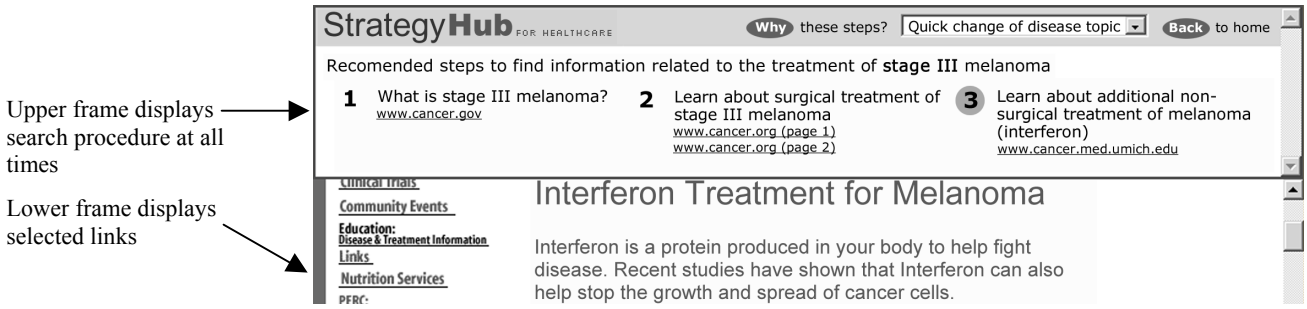

Figure 3. After "melanoma" and "Stage III" are selected in the home page, the Strategy Hub displays the recommended steps of a search procedure in the upper frame, with reliable links for each step. When a link is selected, the associated page is displayed in the lower frame.

in the lower frame. For example, in Figure 3, the link to the University of Michigan site has been selected in the third step, and that page is displayed in the lower frame. This dual frame design is important because it is easy to forget the overall steps in a plan unless it is visible at all times. Furthermore, the circles around the numbers in the steps make explicit which step has a link selected. The dual frame design therefore provides a combination of a context view, which shows you where you are in the procedure, and a focused view of the content. This combination of focus plus context has been found to be critical for interfaces related to search [8].

(3) Provide user control. Although the steps in the recommended search procedure have an inherent order, the interface provides full control to the user who can choose the order in which to visit the steps. This flexibility is essential for two reasons. First, a user might already know the information at a particular step, and therefore find it redundant to visit. Second, a user might be using the Strategy Hub for a specific rather than a comprehensive answer and be satisfied to visit a single step to answer a specific question.

In addition to the above three principles that we felt were critical, the Strategy Hub interface also satisfied all the design desiderata defined by Shneiderman et al. [14]. The interface design is consistent in the way it always provides steps for each topic in the upper frame of the dual-frame window; the quick change of disease topic shown in Figure 3 provides a shortcut to change the disease topic without having to return to the home page; feedback of which topic was selected is displayed above the steps in the dual-frame window; the users know they have searched for all the relevant information because the topic tree is broad and shallow, and the numbered steps provide a sense of closure when users reach the last step; errors are prevented for all selections by graying out wrong selections; reversal of actions is supported by different methods of going back such as the "back to home" button shown in Figure 3; the steps being visible at all times reduces short-term memory loads related to knowing where the user is in the overall search procedure.

To better understand the advantages of providing search procedures via the Strategy Hub, and to identify problems with the interface design, we conducted a pilot study to compare the Strategy Hub with conventional search tools.

PILOT STUDY COMPARING THE STRATEGY HUB TO CONVENTIONAL SEARCH TOOLS

The pilot experiment to test the Strategy Hub was designed with two goals: (1) to analyze the efficiency, effectiveness, and satisfaction of Strategy Hub users, when compared to users of conventional search tools, and (2) to identify problems related with the usability of the interface. We hypothesized that users with two or more years of search experience who used the Strategy Hub would be more effective, efficient, and satisfied in retrieving complete and accurate information from the Web for comprehensive questions, compared to similar users who use conventional search tools. We chose to focus our pilot on expert searchers to establish a baseline comparison for future studies that will focus on novice searchers.

Experimental design. The hypothesis was tested in a 3 X 3 between subject design as shown in Figure 4. The first factor varied across three tool conditions. (1) Any tool or set of tools selected by the user. This represents a realistic condition of how most people currently search the Web for healthcare information [9]. (2) The melanoma page in MEDLINEplus that provides links to reliable sources of melanoma information on the Web. This represents an ideal condition similar to an expert who knows how to navigate MEDLINEplus to reach the melanoma page. (3) The home page of the Strategy Hub for healthcare as shown in Figure 2. This represents a system that provides search procedures in addition to reliable links.

As shown in Figure 4, the second experimental factor varied across three task types: (1) simple melanoma question, (2) difficult melanoma question, and (3) comprehensive melanoma question. The simple and difficult questions were identified in a separate study where two graduate students answered 35 real-world questions related to melanoma using any search method. The answers were ranked by a difficulty score consisting of an aggregate of accuracy and completion (as judged by an expert), and time. The ranking helped identify the simplest and most difficult question shown in Figure 4. The comprehensive question was selected by the skin cancer physicians based on the typical broad treatment question heard in the clinic. Each user was given only one question to control for order

	Any tool chosen by user	Melanoma page in MEDLINEplus	Strategy Hub
Simple question What is a melanoma?	2 users	2 users	2 users
Difficult question Which areas do doctors usually examine when they see patients with melanoma?	3 users	3 users	3 users
Comprehensive question This morning your close friend was diagnosed with melanoma and she is currently undergoing extensive tests to determine her condition. Please try and learn as much as you can about melanoma treatment so you can explain important points related to melanoma treatment to your friend	3 users	3 users	3 users

Figure 4. The pilot experiment had a 3X3 between subject design that varied by tool and task type. Users in each tool condition were given only one task to guard against order effects.

effects because users typically visit many pages during a search and may therefore inadvertently find answers to a later question.

Subjects. Graduate students from the School of Information at the University of Michigan were recruited to take part in the study in return for $25. The recruitment yielded 24 students. All the students had attended a course in search and retrieval, and had two or more years of search experience on the Web. None had searched for information related to melanoma.

Method. The 24 users were randomly distributed in the nine cells as shown in Figure 4. Users in a particular condition were asked to watch a digital video that provided instructions to perform the task, and to answer one question provided in a Word document on their computer. The users were instructed to search for answers using only the approach in their condition, within a maximum of 25 minutes. They were also told to cut and paste into the Word document the paragraphs from the Web that they felt contained the answer, and then to type out their final answer in their own words. This was done to distinguish the retrieval of information, from the interpretation and construction of the final answer.

Regardless of completion, users were asked to stop searching and writing after 25 minutes (this is the upper end of time limits for most studies on the retrieval of information from the Web). A screen capture tool was used to record the interactions on the screen, and a Web logger developed by PARC was used to record the time and occurrence of keystrokes related to clicking a link, scrolling, and using the back button. After they completed the task, the users were instructed to rate how satisfied they were with their search on a 5-point Likert scale. The 8 users in the Strategy Hub condition were also asked to provide written feedback on the interface design.

Analysis. The pilot data were analyzed to quantitatively measure efficiency, effectiveness, and satisfaction, and to qualitatively evaluate feedback about the interface design. Efficiency was measured by (1) total search time (total task time - (cut and paste time + writing time)), and (2) proportion of productive pages (pages where content was cut and paste/total pages visited). Effectiveness was measured by an accuracy score determined by the presence of correct statements, provided by the physicians, in the answer. For example, each comprehensive answer was judged based on whether it described melanoma treatment as stage-based, had the correct treatments for each stage of the cancer, and mentioned clinical trials for advanced stages of the melanoma. An inter-rater reliability test of the accuracy of the answers is in progress. Satisfaction was measured by the responses of how satisfied users were with their search on a 5-point Likert scale.

Results. Only descriptive statistics will be reported given the small number of users in each cell of the pilot design. One user in the Strategy Hub difficult question, and one user in the MEDLINEplus comprehensive question misunderstood the task and provided descriptions of how they searched for the results rather than providing an answer, and therefore were removed from the analysis.

Effectiveness, Efficiency, and Satisfaction. Figure 5 shows the mean measures for effectiveness, efficiency, and satisfaction. As hypothesized, the mean accuracy of the answers for the easy question is 100% across all the conditions. However, users in the Strategy Hub condition have more accurate answers for the difficult and the comprehensive questions. Furthermore, this increase in accuracy does not come at the expense of time; Strategy Hub users have the lowest time for both the difficult and comprehensive questions. This lower time appears to be explained by the comparatively higher percentage of productive pages that the Strategy Hub users visited. Finally, mean satisfaction for the Strategy Hub is highest only for the comprehensive question suggesting that the Strategy Hub may be overkill for specific questions, and better suited for comprehensive questions.

The results therefore suggest that Strategy Hub users are more effective for difficult and comprehensive questions, and in general visit fewer pages that are not productive. It is pertinent to emphasize that the users in this study were skilled at searching, and users in the MEDLINEplus condition were given the melanoma page, not the homepage from where they would have had to search for the melanoma page. This experiment therefore compared the Strategy Hub to other tool conditions in their ideal and expert states. Despite this, the results show an improvement in the search efficiency and effectiveness for difficult and comprehensive questions.

	Simple question			Difficult question			Comprehensive question		
	Any tool Mean (SD)	MEDLINE plus Mean (SD)	Strategy Hub Mean (SD)	Any tool Mean (SD)	MEDLINE plus Mean (SD)	Strategy Hub Mean (SD)	Any tool Mean (SD)	MEDLINE plus Mean (SD)	Strategy Hub Mean (SD)
Efficiency									
Mean search time (in minutes)	7:56 (0:04)	14:34 (3:47)	12:38 (2:39)	11:44 (4:16)	15:43 (2:10)	6:24 (0:35)	13:55 (0:49)	16:13 (1:04)	13:24 (1:22)
Mean percentage of productive pages	29% (30%)	16% (3%)	67% (24%)	14% (9%)	13% (3%)	28% (8%)	22% (8%)	31% (17%)	46% (30%)
Effectiveness									
Mean accuracy	100% (0%)	100% (0%)	100% (0%)	33% (33%)	44% (38%)	66% (0%)	10% (7%)	24% (20%)	26% (17%)
Satisfaction									
Mean satisfaction	5 (0)	4 (0)	4 (0)	4.67 (0.58)	4.33 (1.15)	4 (0)	4 (1)	3 (1.41)	4.67 (0.58)

Any tool = any tool chosen by user; **MEDLINEplus** = melanoma page in MEDLINEplus; **Strategy Hub** = Strategy Hub for healthcare home page

Figure 5. Results from pilot study showing measures for efficiency, effectiveness, and satisfaction for the different questions and conditions in the experiment. Mean percentage productive pages, and mean accuracy, are rounded to the nearest percent. Satisfaction was measured on a 5-point Likert scale where 1=extremely unsatisfied, and 5=extremely satisfied.

Interface design. An analysis of the open-ended feedback provided by the users in the Strategy Hub condition revealed no problems that prevented any user from completing their task. This confirmed the mean satisfaction score of four and above by the Strategy Hub users shown in Figure 5. However, they provided the following five suggestions: allow dual-frame to be turned off, display URL of page being visited at all times, put .gov and .org sites above .com sites, allow brief explanations to be cut and pasted instead of just being roll-overs, provide topic taxonomy based on symptoms and area of body instead of just disease.

While the above self-reported feedback about the interface did not reveal any problems that prevented users from completing the tasks, we analyzed the screen capture videos to probe if the eight Strategy Hub users mapped to the correct nodes in the taxonomy. The analysis revealed that all but one user (P-40) mapped correctly to the nodes best suited to answering a question (i.e. the *Terminology* node for the simple question, *Diagnosis* node for the difficult question, and *Treatment* node for the comprehensive question). P-40, while answering the comprehensive question, first mapped to *Terminology->Defnitions*, and then selected a link suggested by the Strategy Hub containing general information about melanoma. After reading general information about melanoma, he selected a link named *Treatment* within that page leading to a page that was not recommended by the physicians. He therefore never visited the *Treatment* node in the Strategy Hub taxonomy, and hence never had access to a search procedure for treatment. This user had the lowest accuracy score (10%) compared to the other two Strategy Hub users doing this question that did map to the correct node (22%, 44%). This large difference in accuracy motivated us to explore how to address the issue of users never accessing appropriate search procedures suggested by the Strategy Hub.

Implications for Interface Design. Clearly Strategy Hub users are free to navigate right out of the Strategy Hub by following links that were not selected by the experts. Therefore we were initially puzzled how to encourage users like P-40 to stick with the Strategy Hub. An informal study with three novice searchers has shed some light on what might be happening with such users, and a possible solution. The novices performing the comprehensive question (and who also exhibited similar behavior), stated that they thought that a page displayed in the lower frame of the dual frame design was *part of* the Strategy Hub, not a page *pointed to* by the Strategy Hub. This led to the hypothesis that the dual frame design, at least for some users, was not providing enough distinction of what was part, and what was not part, of the Strategy Hub. Furthermore, the dual frame design also presented a technical problem. Some sites are designed such that they take over the frame of the site that linked to it. For example, if a Strategy Hub user clicks on some of the links within the Cancer.org site, the selected page *takes over* the entire browser window, obliterating the Strategy Hub. Users may therefore lose context and never return to the search procedure provided by the Strategy Hub.

The above visual and technical problems with the dual frame design motivated us to redesign the interface to provide a stronger separation between the Strategy Hub and the linked pages. Our new design has no frames, and the links for each step in the search procedure open in a new window. This visual separation of the search procedures provided by the Strategy Hub, and the pages that they point to, addresses the incorrect interpretation that the linked pages are content provided by the Strategy Hub. Furthermore, it eliminates the frame take-over problem described earlier. It is important to note that the shift from a dual frame to a dual window design still allows the user to always have the entire search procedure visible in addition to the linked pages. The design therefore still maintains the focus plus context principle discussed earlier.

In response to the interface feedback, the new design makes URLs visible at all times, and we are also exploring how text from the *Brief explanation* box could be cut and pasted. Finally, we are in consultation with the doctors to explore the utility of alternate taxonomies for topic selection other than what we had provided. Future studies will reveal whether these new design features will lead to more effective interactions with the Strategy Hub.

SUMMARY AND CONCLUSION

Our research was motivated by two observations. (1) Expert searchers have acquired effective and efficient search procedures that guide them to retrieve comprehensive information about a topic from different sources. Such search procedures are necessary given the wide variability of detail and specialization of information on the Web. (2) Novices, searching in unfamiliar domains find it difficult to infer such search procedures from conventional search tools, often leading to the retrieval of incomplete information.

To address the above situation, we collaborated with healthcare search experts to systematically identify search procedures to find comprehensive information for a specific disease. Analysis of the search procedures showed that they indeed did exploit the variability in detail and specialization of information on the Web. Furthermore, the search procedures could be generalized into templates to assist in the identification of search procedures in other domains. We then showed how the search procedures could be made available on the Web through a new form of domain portal called a Strategy Hub using principles of user-centered design.

A pilot study suggested that the Strategy Hub could improve the efficiency, effectiveness and satisfaction of even expert users when attempting to answer comprehensive questions. The pilot study also helped to identify problems in the interface, which motivated us to make changes in the design. Our current research tests the new design on a larger population of novice searchers using the lessons we learned from the pilot experiment.

Although we have shown how search procedures can be provided in a new form of domain portal, we believe the notion of providing such procedural knowledge is much more general. Search procedures can be useful within any large site where there does not exist a one-to-one mapping between a task and a page. In such cases, the retrieval of information from the site would require the user to infer which pages to visit in which order to get a comprehensive understanding of the topic. This can be time-consuming and error-prone. In such situations, search procedures like those we have described, could guide users to appropriate pages in the right order leading to more comprehensive searches. Furthermore, we believe that the automatic categorization of links [e.g. 7], as implemented by commercial search engines such as Vivisimo, could also provide search procedures to guide users to find more comprehensive information.

The notion of providing search procedures, their generalization, and the interface design related to how to provide them are therefore the important contributions in this paper. Besides providing a new direction in the research for search interfaces, search procedures should lead users to be more effective, efficient, and satisfied when finding comprehensive information in unfamiliar domains.

ACKNOWLEDGEMENTS

This research was supported in part by the National Science Foundation, Award# EIA-9812607. The views and conclusions contained in this document should not be interpreted as representing the official policies, either expressed or implied, of NSF or the U. S. Government. The author thanks M. Bates, V. Rajamanickam, and F. Reif, for their contributions.

REFERENCES

1. Bhavnani, S.K. Important components of domain-specific search knowledge. *Proceedings of TREC' 01* (Gaithersburg MD, November 2001), NIST, 571-578.
2. Bhavnani, S.K. Domain-specific search strategies for the effective retrieval of healthcare and shopping information. *Proceedings of the CHI'02* (Minneapolis MN, April 2002), ACM Press, 610-611.
3. Bhavnani, S.K., Bichakjian, C.K., Schwartz, J.L., Strecher, V.J., Dunn, R.L., Johnson, T.M., & Lu, X. Getting patients to the right healthcare sources: From real-world questions to Strategy Hubs. *Proceedings of AMIA'02* (San Antonio TX, November 2001) 51-55.
4. 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 20*, 1 (2002), 134-141.
5. Chi, E., Pirolli, P., & Pitkow, J. The scent of a site: A system for analyzing and predicting information scent, usage, and usability of a Web site. *Proceedings of the CHI' 00*, (The Hague, Netherlands, April 2000), ACM Press, 161-168.
6. Crowston, K., & Williams, M. Reproduced and emergent genres of communication on the World-Wide Web. *Proceedings of HICSS–30* (Maui HI, January 1997).
7. Dumais, S., Cutrell, E., & Chen, H. Optimizing search by showing results in context. *Proceedings of the CHI' 01* (Seattle WA, April 2001), ACM Press, 277-284.
8. Egan, D.E., Remde, J.R., Landauer, T.K., Lochbaum, C.C., & Gomez, L.M. (1989). Behavioral evaluation and analysis of a hypertext browser. *Proceedings of the CHI'89* (Austin TX, May 1989), ACM Press, 205-210.
9. 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. Avail: http://www.pewinternet.org/reports/toc.asp?Report=26 (2000).
10. Furnas, G., Landauer, T., Gomez, L., & Dumais, S. The vocabulary problem in human-system communication. *Communications of the ACM 30*, 11 (1987), 964-971.
11. Kirk, T. Problems in library instruction in four-year colleges. In: Lubans, John, Jr. (ed.), *Educating the library user*, 83-103. R. R. Bowker, New York NY, 1974.
12. Pratt, W., Hearst, M., & Fagan, L. A Knowledge-Based Approach to Organizing Retrieved Documents. *Proceedings of the AAAI '99* (Orlando FL, July 1989), AAAI Press, 80-85.
13. Shute, S., & Smith, P. Knowledge-based search tactics. *Information Processing & Management 29*, 1 (1993), 29-45.
14. Shneiderman, B., Byrd, D., & Croft, W.B. Clarifying search: A user interface framework for text searches. *DLIB Mag 3*, 1 (1997).
15. Thottam, J. Search smarter. *On Magazine*, November 2001, 33-37.
16. Voss, D. Better searching through science. *Science 293* (2001), 2024-2026.