

Lessons Learned from Building the iMed Intelligent Medical Search Engine

Gang Luo

IBM T.J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532, USA

luog@us.ibm.com

Abstract — Searching for medical information on the Web has become highly popular, but it remains a challenging task because searchers are often uncertain about their exact medical situations and unfamiliar with medical terminology. To address this challenge, we have built an intelligent medical Web search engine called iMed. iMed introduces and extends expert system technology into the search engine domain. It uses medical knowledge and an interactive questionnaire to help searchers form queries. This paper reports the lessons we learned from building the iMed system. We believe that many of these lessons can be applied to other medical search engines as well. We systematically discuss important issues in the new field consumer-centric intelligent medical search, including input interface, output interface, search system, medical knowledge base, help system, and testing.

I. INTRODUCTION

Today, ordinary Internet users are increasingly using Web search engines (WSEs) to search for medical information on the Web (6% of American Internet users on an average day [13]). As the healthcare industry is moving toward a more consumer-centric focus, this trend is expected to last in the foreseeable near future [7].

Since October 2005, several medical WSEs have been launched, including Healthline [5] and Google Health [3]. These systems have some features tailored to medical search. For a query Q , they can recognize the medical phrase P in Q and automatically add P 's synonyms into Q . Moreover, to avoid the disturbance of low-quality pages from irrelevant Web sites in the search results, they use a vertical WSE approach that crawls Web pages from a few selected, high-quality medical Web sites rather than all the Web sites.

Nevertheless, medical search has its own unique requirements that distinguish itself from traditional Web search. Existing medical WSEs have not sufficiently addressed these unique requirements and cannot satisfactorily meet the needs of medical information searchers. For example, all existing medical WSEs assume that searchers can form appropriate queries by themselves. In many cases, the medical information searcher is uncertain about the problem he is facing and unaware of the related medical terminology (e.g., panophthalmitis). As a result, it is often difficult for him to choose a few accurate medical phrases as a starting point for his search.

To address the limitations of existing systems, we have built an intelligent *medical* WSE (IMWSE) called iMed whose individual components are described in [8], [9], [10], [11]. iMed introduces and extends expert system technology into the WSE domain. It extensively uses medical knowledge

and novel user interfaces including interactive questionnaire to help ordinary Internet users form queries and search for medical information. During its design and implementation, we learned many valuable lessons and identified a few important issues related to intelligent medical search. Since consumer-centric intelligent medical search is a new field, we describe these lessons and issues in this paper and hope they can benefit other people working in the medical search area. We also present a whole system view of iMed together with the design rationale.

The rest of the paper is organized as follows. Section II presents an overview of iMed. Section III describes our general principles of designing, implementing, and testing an IMWSE. Section IV shows some experimental results.

II. OVERVIEW OF IMED

In this section, we provide an overview of iMed. Our presentation focuses on the underlying reasoning process of designing, implementing, and testing iMed. The details of our implementation techniques are described in [8], [9], [11]. In general, the unique requirements of medical search affect every aspect of the design of an IMWSE. The combined effect of these requirements is so large that IMWSEs and traditional WSEs have radically different architectures. In the following, we first describe the architectural difference. The unique requirements of medical search are mentioned when we present the individual components of an IMWSE.

A. System Architecture

Fig. 1 shows the architecture of a traditional WSE. It has three components: the input interface, the search system, and the output interface. These three components can be arranged in a linear order and interactions occur only between neighboring components in a unidirectional way. Users can see both the input interface and the output interface but not the search system.




Fig. 1. Architecture of a traditional search engine.

Fig. 2 shows the architecture of an IMWSE. Compared to the architecture of a traditional WSE, it has two more components: the help system and the medical knowledge base. Interactions occur between almost each pair of components whereas most interactions are bidirectional. Users can see the input interface, the output interface, and the help system, whereas they cannot see the medical

knowledge base and the search system. Next, we discuss the individual components of an IMWSE one by one.




Fig. 2. Architecture of an intelligent medical search engine.

B. Input Interface

When users use an IMWSE, they first see the input interface. An ordinary user typically has little medical knowledge and is uncertain about the problem he is facing. Consequently, he often cannot well describe his medical situation using a few keywords. One way to address this problem is to remove existing WSEs' query length limits [12] and to allow users to input long queries, describing their symptoms and situations in detail in plain English. The medical WSE takes the burden of identifying important keywords in a query and uses these keywords to perform search. This is the approach adopted in our first prototype medical WSE MedSearch [12].

Nevertheless, medical cases are often complicated. Frequently an ordinary user needs several paragraphs to describe his situation, whereas a large part of his description is unnecessary because he has no idea about what information is important for diagnosing his medical problem. Due to the same reason, the user often forgets to mention some important information. Moreover, the medical WSE has limited capability of using existing natural language processing techniques to identify important keywords in a query. All these undesirable effects significantly deteriorate the quality of query keywords and hence the quality of search results.

To offer the greatest convenience to users, an IMWSE should provide two kinds of input interfaces simultaneously. First, if the user knows the appropriate query keywords (e.g., the exact name of the disease), he can use the traditional keyword search interface to find desirable search results. Second, when this is not the case, the WSE should provide a guided interface to help the user input important information about his medical situation. In the rest of the paper, we focus on the guided input interface.

During a doctor's office visit, the doctor collects relevant medical information about a patient by asking questions. Motivated by this observation, our guided input interface is based on questionnaires [8] and mimics the way that doctors interact with patients. As shown in Fig. 3 and Fig. 4, the user first selects one or more symptoms and signs from a list of symptoms and signs. Then the IMWSE asks questions related to these selected symptoms and signs. In the case that all the important information about the user's medical situation is covered by the questionnaire, the user provides information by selecting symptoms, signs, and question answers without performing manual typing.

In theory, to maximize its coverage of all possible medical cases, the questionnaire could include as many symptoms, signs, and questions as the entire human medical knowledge permits. However, this would make the questionnaire overly complicated. In general, users prefer using simple interfaces and finding desirable search results as soon as possible. Hence, we need to control the questionnaire's complexity by focusing on common medical problems [8]. The questionnaire needs to be well organized so that most users can quickly find those items corresponding to their cases. Backup options (e.g., backup questions) need to be provided when a user does not know the answer to a question. At any time, a user can stop the input process and obtain search results based on his inputs so far. Also, the questionnaire needs to be clearly written so that users can easily make selections without either becoming confused or having to laboriously consult other resources for clarification.

Symptoms and Signs	
<input checked="" type="checkbox"/> Abdominal Pain	<input type="checkbox"/> Irregular Periods
<input type="checkbox"/> Backache	<input type="checkbox"/> Painful Periods
<input type="checkbox"/> Belching, Bloating and Flatulence	<input type="checkbox"/> Nausea and/or Vomiting without Abdominal Pain
<input type="checkbox"/> Breast Lumps	<input type="checkbox"/> Foot Pain
<input type="checkbox"/> Chest Pain	<input type="checkbox"/> Leg Pain
<input type="checkbox"/> Colds, Flu and Stuffy Nose	<input type="checkbox"/> Arm Pain
<input type="checkbox"/> Constipation	<input type="checkbox"/> Abnormal Heartbeat Sensation
<input type="checkbox"/> Cough	<input type="checkbox"/> Shortness of Breath
<input type="checkbox"/> Diarrhea	<input type="checkbox"/> Skin Problems
<input type="checkbox"/> Dizziness/Light-headedness and Vertigo	<input type="checkbox"/> Sore Throat
<input type="checkbox"/> Earache	<input type="checkbox"/> Swollen Legs
<input type="checkbox"/> Facial Pain	<input type="checkbox"/> Urethral Discharge and Painful Urination
<input type="checkbox"/> Fatigue	<input type="checkbox"/> Vaginal Discharge and Itching
<input type="checkbox"/> Fever	<input type="checkbox"/> Vision Problems
<input type="checkbox"/> Forgetfulness	<input type="checkbox"/> Voiding Disorders and Incontinence
<input type="checkbox"/> Headache	<input type="checkbox"/> Weight Gain and Weight Loss
<input type="checkbox"/> Heartburn and Indigestion	<input type="checkbox"/> Others
<input type="checkbox"/> Insomnia	

Previous Next

Fig. 3. The first level of the questionnaire.

Abdominal (stomach) pain

How long have you had abdominal pain?

I have had it for less than 3 months.

I have had it at least 3 times over 3 months or more.

No answer.

other inputs:

Previous Next Finish

Fig. 4. The first question page that is generated for the symptom "abdominal pain."

To make the questionnaire clear and easy to use, we adopt the following principles. The questionnaire uses layman terms whenever possible, e.g., to explain each used, difficult medical concept. Since a single medical concept can have multiple layman names known to different people, all these names are listed simultaneously when that medical concept is mentioned to facilitate recognition. For instance, the medical concept of libido is written as "libido (sexual desire, sex drive)." When asking whether a medical condition of the

user is chronic or acute, the distinction between chronic and acute is clearly marked as a specific number of weeks or months that this medical condition has lasted. When asking about a medical concept representing a category of items, we also list those items if possible so that the user can quickly know whether his case falls into that category. Specifically, we list all the generic names and brand names of the drugs in a drug category. For instance, the question answer “I take beta blockers.” is written as “I take beta blockers, such as acebutolol (Sectral), atenolol (Tenormin), ...” Also, we list all the symptoms and signs in a symptom category. For example, the question “Do you have any focal neurologic signs?” is written as “Do you have any focal neurologic signs, such as personality change, paralysis, ... ?”

C. Medical Knowledge Base

The extensive use of medical knowledge is the key to making a medical WSE intelligent and to quickly obtaining high-quality medical search results. An IMWSE stores in its medical knowledge base all its medical knowledge, which is used to support its input interface, output interface, help system, and search system.

For example, in the questionnaire-based input interface, all the symptoms, signs, and questions are stored in the medical knowledge base in the form of diagnostic decision trees. These decision trees are described in a few medical textbooks [1], [2], [4], [6] and can be used to find the names of potential diseases corresponding to a user’s medical situation. One such decision tree is shown in Fig. 5.

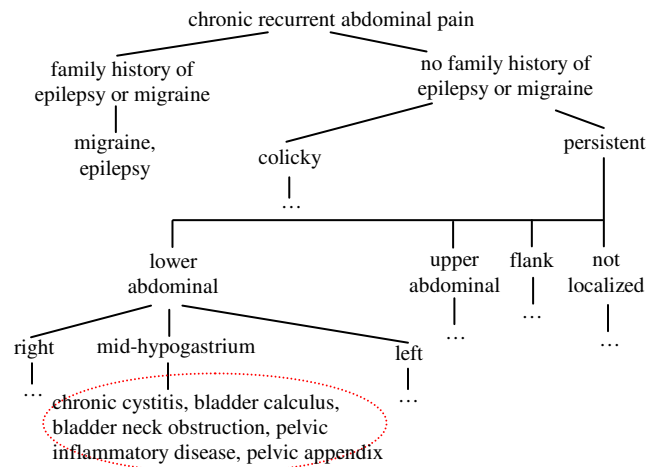


Fig. 5. The diagnostic decision tree for the symptom “chronic recurrent abdominal pain.”

In general, for a given symptom or sign, multiple diagnostic decision trees exist in different medical textbooks. Some of these decision trees have small depths whereas the others have large depths. Ideally these decision trees should be combined into a single one with a small depth. This combined diagnostic decision tree is stored in the knowledge base. In this way, a user only needs to answer a few questions, rather than many questions, before he can obtain search results. Also, backup options can be easily provided when a user does not know the answer to a question.

The knowledge base stores the diagnostic decision trees for not-extremely-rare medical problems and other medical knowledge (e.g., disease incidence rates) most needed by the IMWSE [11]. It does not store the most extremely complex part of human medical knowledge. Otherwise the search reasoning process will become overly complicated and unmanageable. At the same time, users will become overloaded with a large amount of information that is unlikely to be useful.

In principle, the medical domain is a closed one. Much functionality that is difficult or impossible to provide in open domain search can be offered in intelligent medical search by using medical knowledge appropriately. For example, consider the task of suggesting related phrases to help the user quickly digest search results and refine his inputs [8], [12]. In the medical domain, medical phrases are of a limited number whereas much knowledge is available about their relationships. In contrast, phrases in the open domain are almost of an infinite number and can have complex relationships depending on the context. Compared to the case of open domain search, this task becomes much easier in intelligent medical search because of the reduced search space and available medical knowledge. Nevertheless, in order to know what kind of medical knowledge to use and how to use it properly, we often need to consult medical references and collaborate with medical professionals.

D. Output Interface

The search results are displayed in the output interface. Typically, there are multiple topics (e.g., diseases) potentially relevant to the user’s medical situation. For each such topic, the user can be interested in several aspects (e.g., symptom, diagnosis, and treatment). To satisfy the different needs of various users simultaneously, the output interface needs to show the search results of all the potentially relevant topics and their most common aspects. Since a typical user has little medical knowledge and usually does not remember all his desirable aspects of a topic, the IMWSE should automatically provide these aspects using its medical knowledge stored in the medical knowledge base.

The IMWSE can retrieve many search results for a user’s medical situation. If all these search results are presented in a traditional sequential order, their internal structures and relationships will be lost. This brings the user much inconvenience. For example, the user can neither quickly find those search results on a relevant topic nor easily skip the search results on an irrelevant topic.




Fig. 6. The hierarchical structure of the output interface.

As shown in Fig. 6, to provide the greatest convenience to users, all the search results should be structured into a multi-

level hierarchy that has explicitly marked medical meanings [9]. Also, overview Web pages need to be provided at the top levels of the hierarchy. In this way, users can efficiently navigate among the search results and quickly obtain the desired information.

In general, a medical information searcher plays different roles at various times. For example, at the beginning of the search process, he usually prefers to scan all the topics potentially relevant to his medical situation. After identifying the topics most important to him, he often would like to check these topics in detail. At present, it is still an open problem to identify how many common roles medical information searchers usually play. Nevertheless, an IMWSE should realize that different roles have various needs. In the case that one size does not fit all, the input interface, the output interface, and the help functionality can be specifically designed for each individual role.

E. Help System

Since most users have little medical knowledge, they frequently encounter various kinds of problems and need much help during the entire medical search process. Many such problems are specific to ordinary people and are not problems for medical professionals. Identifying these problems requires repeatedly performing user study with ordinary people instead of with medical professionals.

The help system of the IMWSE provides different kinds of suggestions to facilitate the search process [8], [11], such as (1) suggesting diversified, related medical phrases to help the user quickly digest search results and refine his inputs, (2) suggesting symptoms and signs related to the user's medical situation, and (3) suggesting alternative answers to the questions asked by the WSE in case that the user answers questions incorrectly. The help system also provides explanations of symptoms, signs, asked questions, and suggested medical phrases in layman terms. All the suggestions and explanations need to be well organized in the user interface so that users can efficiently use them. This should be done carefully, otherwise users can easily be overwhelmed with the large number of suggestions and explanations. In some cases, new graphical-based user interface techniques need to be developed in order to properly display these suggestions and explanations [11].

F. Search System

In traditional WSEs, users are responsible for forming queries. This rule needs to be changed in medical search, as medical information searchers often have difficulty in forming queries due to lack of medical knowledge. Actually an important feature of an IMWSE is that its search system can automatically form queries. These queries are obtained using both the information the user provides in the questionnaire-based input interface and the medical knowledge stored in the knowledge base [8]. For example, the user's selected symptoms, signs, question answers can be converted into disease names using diagnostic decision trees. A query can include the disease keywords, the symptom (or sign) keywords, and the question answer keywords.

In a traditional WSE, the user inputs one query at a time. In contrast, the search system of an IMWSE forms multiple queries simultaneously. Each query is a combination of a topic (e.g., a disease) and an aspect. The search results of all these queries are combined together and returned to the user. This automatic query formation is possible because the medical domain is a closed one. There, the names of the interesting topics and aspects are known in advance.

G. Testing

Testing is indispensable for discovering problems and obtaining useful feedback on improving the IMWSE. To maximize the benefits obtained from testing, the testing scenarios should closely reflect real usage scenarios. More specifically, most users of an IMWSE are ordinary people without much medical background rather than medical professionals. Hence, the IMWSE should mainly be tested by ordinary people. Moreover, in the majority of the cases, users use the IMWSE to search for information related to common medical problems. Consequently, most test cases should also be about common medical problems.

During testing, users may not always be able to directly point out the shortcomings of the IMWSE whereas they usually have certain complaints. Their common complaints provide hint to the problems with the IMWSE. By following the complaining users' search and reasoning process and carefully reading the search results, the WSE designer can often find clues on how to improve the IMWSE. This is often a time consuming process that easily takes days or months. The WSE designer needs to be patient. To reduce cost, testing can start with one or two users before moving to a larger-scale user study. Easy problems can often be discovered with a small number of users.

III. GENERAL PRINCIPLES

During the design, implementation, and testing of an IMWSE, we think that the following principles are generally applicable:

1. Whenever possible, automatically offer users what they want instead of waiting until they explicitly ask. Users often do not remember or know exactly what they want due to lack of medical knowledge, whereas they usually can tell whether the presented medical information is helpful. Making selections is much easier for users than coming up with the right search keywords.
2. Users care about their health and are willing to spend time on medical search. Reading search results is much more time consuming than waiting for search results to be generated. Hence, IMWSEs can be a little slower than traditional WSEs if this sacrifice can lead to higher-quality, better organized search results.
3. Focus on common medical situations. These are the cases that most people care about and hence will have the biggest impact on the user group. Keep things simple, stupid, and well organized. Simplicity is more important than completeness. It is not worth to make the system more complex to cover a few difficult medical

situations if this hurts the effectiveness that common medical situations are handled.

4. Users care more about finding good search results than finding all the relevant information on the Web. On the other hand, users need second-opinion information to confirm their impressions. Hence, it is important to display high-quality information from more than one source but unnecessary to present the same information repeatedly from all possible sources.
5. Medical search is generally harder than traditional Web search. Most users are ordinary people without much medical background. They encounter many problems that medical professionals will not run into. It is essential to provide the greatest convenience to users and to offer them various kinds of help throughout the entire search process.
6. User interface is the only thing users see. Everything else is used to support the user interface. Always design the user interface first by focusing on simplifying and facilitating the user's task as much as possible. Use layman terms in the user interface whenever possible.
7. Search is always an iterative procedure. Human involvement may be unavoidable during the search process, but the less, the better. Use medical knowledge throughout the entire search process whenever possible, but automatically. Ordinary users are unlikely to have deep medical knowledge and hence their involvement should not be based on the assumption of possessing such knowledge.
8. Obtain feedback through testing, mainly using ordinary people and common medical cases that are most critical to the success of the system. Building an IMWSE requires interdisciplinary knowledge. Check medical references. Talk with both ordinary people and medical professionals to figure out what is practically important for ordinary people, medically meaningful, and feasible by utilizing the existing human medical knowledge.
9. The medical domain is a closed one. By limiting the search space and utilizing medical knowledge, things that are difficult (or infeasible) to do in open domain search can become relatively easy (or possible) in medical search. Keep the difference between medical search and traditional Web search in mind rather than overly trusting the experience gained from building traditional WSEs.
10. Solve one problem at a time. Do not be overly greedy, especially at the early stage of building an IMWSE. Fast prototyping is important and can quickly identify improper designs. It can also help discover important areas and problems that are difficult to realize by pure thinking. For this purpose, testing one medical case thoroughly is often more helpful than testing multiple medical cases carelessly, whereas obtaining more hands-on experience is always beneficial.
11. IMWSE is not God and cannot solve all the problems. It is used to provide medical information rather than replacing physicians. The expectation on the IMWSE has to be realistic.

V. EXPERIMENTAL RESULTS

In the past, we have evaluated iMed [8], [9], [11] by crawling a large number of medical Web pages from the Internet and by using a wide range of medical scenarios. Our results show that iMed can significantly outperform existing medical WSEs such as Healthline [5] and Google Health [3].

To give the reader a feeling of the contents returned by iMed, we present detailed results of the returned Web pages and the suggested medical phrases for a typical query scenario that corresponds to choosing "no family history of epilepsy or migraine," "persistent," "lower abdominal," and "mid-hypogastrium" for the symptom "chronic recurrent abdominal pain" (see Fig. 5). At the first level of the search result hierarchy, some returned relevant Web pages are shown in Table I. The suggested relevant medical phrases include ulcer (rank 1), pyelonephritis (rank 2), ileitis (rank 3), and pancreatitis (rank 6). In general, for a query scenario Q_s , iMed can find several relevant Web pages and medical phrases describing multiple topics related to Q_s .

TABLE I. SOME RETURNED RELEVANT WEB PAGES.

rank	URL	topic
1	www.webmd.com/urinary-incontinence-oab/interstitial-cystitis-painful-bladder-syndrome	chronic cystitis
2	www.webmd.com/kidney-stones/kidney-stone-analysis	bladder calculus
4	www.webmd.com/a-to-z-guides/pelvic-inflammatory-disease-topic-overview	pelvic inflammatory disease

ACKNOWLEDGMENT

We thank Chunqiang Tang for helpful discussions.

REFERENCES

- [1] *American Medical Association Family Medical Guide*, 4th ed. John Wiley & Sons, 2004.
- [2] R.D. Collins, *Algorithmic Diagnosis of Symptoms and Signs: Cost-Effective Approach*. Lippincott Williams & Wilkins, 2002.
- [3] (2008) Google Health homepage. [Online]. Available: <http://www.google.com/Top/Health>.
- [4] D.R. Goldmann, *American College of Physicians Complete Home Medical Guide*. DK Publishing, 2003.
- [5] (2008) Healthline homepage. [Online]. Available: <http://www.healthline.com>.
- [6] A.L. Komaroff, *Harvard Medical School Family Health Guide*. Free Press, 2004.
- [7] S. Lohr, Google and Microsoft look to change health care. New York Times, August 14, 2007.
- [8] G. Luo, "Design and evaluation of the iMed intelligent medical search engine," in *Proc. ICDE'09*, 2009, pp. 1379-1390.
- [9] G. Luo, "Intelligent output interface for intelligent medical search engine," in *Proc. AAAI'08*, 2008, pp. 1201-1206.
- [10] G. Luo and C. Tang, "Challenging issues in iterative intelligent medical search," in *Proc. ICPR'08*, 2008, pp. 1-4.
- [11] G. Luo and C. Tang, "On iterative intelligent medical search," in *Proc. SIGIR'08*, 2008, pp. 3-10.
- [12] G. Luo, C. Tang, H. Yang, and X. Wei, "MedSearch: a specialized search engine for medical information retrieval," in *Proc. CIKM'08*, 2008, pp. 143-152.
- [13] C. Sherman. (2005) Curing medical information disorder. [Online]. Available: <http://searchenginewatch.com/showPage.html?page=3556491>.