

Challenging Issues in Iterative Intelligent Medical Search

Gang Luo Chunqiang Tang
IBM T.J. Watson Research Center
{luog, ctang}@us.ibm.com

Abstract

Searching for medical information on the Web is highly popular these days. To facilitate ordinary people to perform medical search and preliminary disease self-diagnosis, we have built an intelligent medical Web search engine called iMed. iMed introduces and extends pattern recognition and expert system technology into the search engine domain. It uses medical knowledge and an interactive questionnaire to help searchers form queries. Due to searchers' limited medical knowledge and the task's inherent difficulty, searchers often cannot find desired search results in a single pass and have to search iteratively for multiple passes. For this purpose, iMed provides an iterative search advisor that guides searchers to refine their inputs. Based on our experience in building and using iMed, this paper summarizes the common difficulties faced by ordinary medical information searchers and the research issues that deserve attention from people working in the pattern recognition and medical search areas.

1. Introduction

Today, ordinary Internet users are increasingly using Web search engines to search for medical information on the Web (6% of American Internet users on an average day [12]). Since October 2005, several medical Web search engines have been launched, including Healthline [5], Google Health [3], SearchMedica [13], and Medstory [11]. They use the traditional keyword query interface, which works well when the searcher clearly knows his medical situation. For instance, a searcher knows that he has high cholesterol and wants to learn about appropriate diet for himself. However, 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 [10].

To address this problem, we have built a prototype intelligent medical search engine called iMed [7, 8, 9]. iMed introduces and extends pattern recognition and expert system technology into the search engine domain. It uses medical knowledge and an interactive questionnaire to help searchers form queries, search medical information, and perform preliminary disease self-diagnosis. iMed performs better than existing medical search engines and makes medical search easier than before, while medical search remains as a challenging problem. Even for physicians with much medical experience, performing medical search is often a difficult task [2, 6]. For ordinary Internet users with little medical knowledge, we expect their medical search performance to be even worse.

Frequently, searchers cannot find desired search results in a single pass and have to search iteratively for multiple passes. Since intelligent medical search engines differ significantly from traditional medical search engines, searchers face different difficulties when using intelligent medical search engines. Moreover, ordinary searchers without much medical background frequently encounter many problems that medical professionals typically would not run into.

In the rest of the paper, we first give a brief overview of iMed, and then report the lessons we learned in building and using the iMed system, especially its iterative search advisor. Our focus is on the common difficulties faced by ordinary medical information searchers and the research issues that deserve attention from people working in the pattern recognition and medical search areas.

2. Overview of iMed

iMed leverages its built-in medical knowledge in the form of diagnostic decision trees written by medical professionals [1, 16, 17, 18, 19]. As shown in Figure 1, each diagnostic decision tree corresponds to either a subjective symptom (e.g., fatigue) or an objective sign

(e.g., hypertension). Each non-leaf, non-root node of a diagnostic decision tree corresponds to an answer to a question that iMed can ask. Each medical phrase in the leaf node of a diagnostic decision tree (possibly in combination with the searcher's other keyword inputs) can become a query that iMed uses.

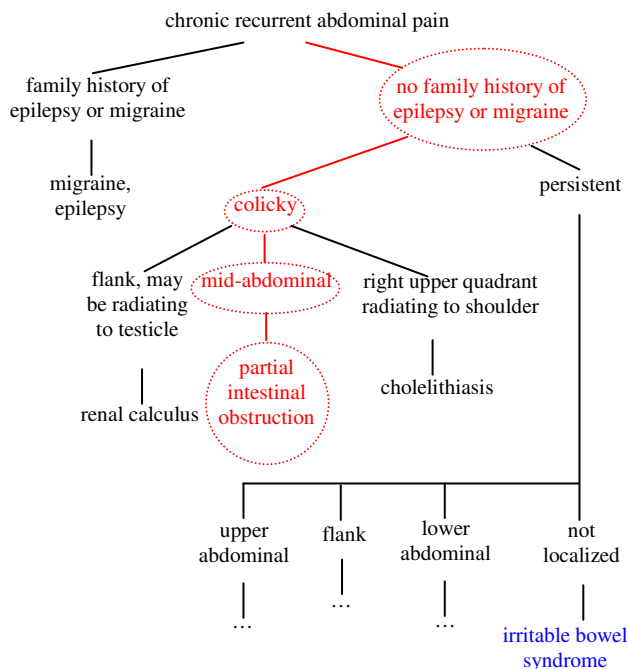


Figure 1. The diagnostic decision tree for the symptom “chronic recurrent abdominal pain.”

iMed uses diagnostic decision trees to help the searcher form queries. The searcher first selects one or more symptoms and signs from a list of 267 symptoms and signs [1]. This list covers most chief complaints with which physicians are confronted. Then iMed asks questions related to these selected symptoms and signs. Based on the searcher's answers to the questions, iMed navigates the corresponding diagnostic decision trees and automatically forms multiple queries. Each query is used to retrieve some related Web pages. iMed's search results include the Web pages retrieved for all these queries [9].

For example, Figure 1 shows the diagnostic decision tree in Collins [1] for the symptom “chronic recurrent abdominal pain.” If “chronic recurrent abdominal pain” is the only symptom chosen by the searcher, iMed's first question is “Is there a family history of epilepsy or migraine?” If the searcher answers “no” to this question, iMed's next question is “Is the pain colicky or persistent?” If the searcher answers “colicky” to the second question, iMed continues to ask “What is the location of the pain?” If the

searcher's answer to the third question is “mid-abdominal,” iMed forms multiple queries including partial intestinal obstruction. (Medical phrases in the non-selected leaf nodes of the diagnostic decision tree also form queries, but with lower weights.) A detailed description of iMed and its use for preliminary disease self-diagnosis is available in [7, 8, 9].

3. Challenges and Lessons Learned

Consumer-centric intelligent medical search is a relatively new field. In this section, we present in detail the lessons we learned from building and using the iMed system, and hope that our experience will be useful for others as well. We believe that many of the lessons we learned are not specific to the iMed system and can be applied to general intelligent medical search. Moreover, new pattern recognition techniques need to be developed to address the challenging issues that ordinary medical information searchers face in performing iterative intelligent medical search.

3.1. Combining Statistical Analysis with Domain Knowledge and User Intelligence

The most important lesson is that the medical search problem cannot be solved using pure information retrieval techniques that largely rely on statistical text analysis, as these techniques can neither well understand the deep semantics of searchers' intents nor well utilize the large amount of available medical practice experience. Medical search is special in that it focuses on the relatively closed medical domain, where much medical knowledge has been well documented, e.g., in the form of diagnostic decision trees. Since an ordinary searcher often has difficulty in clearly describing his medical situation, traditional information retrieval techniques frequently cannot effectively process his keyword inputs. Nevertheless, with guided inputs from the searcher in the form of selection choices, the performance of automated algorithms can be significantly boosted.

We think that the best way to practically address the challenges in medical search is to combine medical knowledge with pattern recognition and information retrieval techniques while taking into account human factors. Essentially, a medical search engine needs to make the best out of three factors to maximize its performance: (1) medical experts' domain knowledge, (2) searchers' intelligence, and (3) the processing power of automatic machine analysis techniques. Therefore, it is mandatory for researchers in this field to have a broad background and to take an interdisciplinary approach.

3.2. Symptoms and Signs vs. Question Answers

The number of symptoms and signs covered in iMed's questionnaire is much larger than the number of questions that iMed asks during a user search session. Therefore, choosing proper symptoms and signs is both crucial and generally more difficult than answering questions appropriately. In general, this is a challenging and important research problem that needs continued endeavor and deserves attention from the research community. Below, we describe our experience on this issue.

iMed classifies all the symptoms and signs into multiple categories. Nevertheless, selecting symptoms and signs is still often a tricky task for several reasons. A symptom or sign can be related to multiple categories but it is only shown in one category. It is time-consuming to check all the 267 symptoms and signs covered in iMed's questionnaire. To be worse, many symptoms and signs have difficult medical names (e.g., pneumaturia) and searchers need to check their detailed medical definitions provided by iMed to make selections. Actually, we have seen cases where searchers do not even know which symptoms and signs to start with at the very beginning of the search process. Moreover, when a person is sick, he can have multiple symptoms and signs, and may even feel uncomfortable everywhere. In this case, it is best to start with his chief complaint, i.e., his most important symptom or sign. However, finding chief complaints is a nontrivial task for an ordinary person without rigorous medical training. On the other hand, if the searcher simply selects all the symptoms and signs that seem to be at least marginally relevant, he can easily be swamped by a lot of noise information and cannot find the desired information.

Selecting inappropriate symptoms and signs is generally more detrimental than providing improper answers to the questions. This is because if the searcher chooses appropriate symptoms and signs, the correct disease d will be covered in the corresponding diagnostic decision trees. In this case, if improper answers are provided to the questions asked by iMed, the query that iMed forms about d will have lower weight than that of some other queries formed by iMed. Consequently, the Web pages P about d will be ranked low among all the Web pages returned by iMed. However, if the searcher is patient enough to read many Web pages returned by iMed, he can still find P and thus d . Moreover, since the number of alternative answers is limited, the searcher may find d through multiple trials with iMed's help. In contrary, if the searcher selects inappropriate symptoms and signs, the correct disease d will not be covered in the

corresponding diagnostic decision trees and hence none of the queries formed by iMed will be related to d . As a result, the searcher is unlikely to find any Web page about d irrespective of how many Web pages returned by iMed is read by him.

3.3. Classification of Improper Selections

It is common that medical information searchers make improper selections when choosing symptoms and signs and answering questions. A good understanding of the nature of these improper selections can be helpful to future medical search engine designers. Based on our experience, we classify these improper selections into three categories. Improper selections from every category are common. Therefore, an intelligent medical search engine should be designed to handle all three categories of improper selections rather than being optimized for a specific category. This is another challenging research problem that deserves attention from the research community. Below, we describe our experience with the improper selections that searchers often make.

In the first category, the searcher can realize his inappropriate selections if he has the opportunity to see the correct diagnosis and to read the corresponding Web pages. Next time when the searcher faces a similar situation, he can learn from his past experience (mostly in the form of textbook-style knowledge) and reduce the likelihood of making improper selections. However, such a likelihood can never be reduced to be zero. This can be illustrated by an analogy to people taking exams. The more exams a person has practiced, the better he will perform in future exams. Nevertheless, it is unrealistic to expect everybody who is well prepared to obtain perfect scores in all the exams.

In the second category, the searcher can roughly realize his inappropriate selections if he has the opportunity to see the correct diagnosis and to read the corresponding Web pages. However, next time when the searcher faces a similar situation, he may still make improper selections. Such cases are common in practice, as medical situations vary case by case, and a gap exists between textbook knowledge and medical practice. As an analogy, in order to obtain his license to practice medicine, every medical student has to go through a lengthy internship process to obtain essential hands-on experience. Actually, without such an internship process, even straight-A students from the best medical schools can easily lose direction when facing real world medical problems [4].

In the third category, the searcher cannot realize his inappropriate selections even if he can see the correct diagnosis and read the corresponding Web pages. Such

cases are not unusual, as many medical situations are inherently fuzzy and even experienced medical professional can become confused and make wrong diagnoses. In fact, according to several studies, doctors' misdiagnosis rates are often above 20% [4].

4. Iterative Search Advisor and Open Issues

Recently, we developed an *iterative search advisor* in iMed [8] to address the challenges described in Section 3. This advisor integrates medical and linguistic knowledge to help searchers improve search results through iterative search. It helps the searcher in the following ways. First, relevant symptoms and signs are automatically suggested based on the searcher's description of his situation. Second, instead of taking for granted the searcher's answers to the questions, iMed ranks and recommends alternative answers according to their likelihoods of being the correct answers.

With a proper iterative search advisor, we expect iMed to work more effectively than medical expert systems [14, 15], as iMed allows an iterative search process and gives the searcher multiple chances while medical expert systems usually only give the user a single chance. Nevertheless, the iterative search advisor only alleviates rather than eliminates the common difficulties faced by ordinary medical information searchers. Moreover, the iterative search advisor does not make the returned search result Web pages easier to understand, while searchers frequently spend hours on laboriously reading and rereading these Web pages that are full of unfamiliar medical terminologies. These are areas for future research.

5. Conclusion

This paper presents a few challenging issues in iterative intelligent medical search. Consumer-centric intelligent medical search is a relatively new field and many problems remain open there. Intelligent medical search engines still need much improvement to provide the greatest convenience to ordinary medical information searchers, while we have already seen some promising results showing that intelligent medical search engines frequently outperform traditional medical search engines. We note that addressing the challenging issues in intelligent medical search requires interdisciplinary knowledge of pattern recognition, expert system, and Web search. An iterative search advisor that combines knowledge from multiple domains is a key technology in addressing these challenges.

6. Acknowledgements

We thank Haiyan Chen, Curt J. Ellmann, Leiguang Gong, Jiuxing Liu, Linda Schumer, Selena Thomas, Ying-li Tian, Jing Wang, Leiping Wang, and Hong Xu for helpful discussions.

References

- [1] R.D. Collins. Algorithmic Diagnosis of Symptoms and Signs: Cost-Effective Approach. Lippincott Williams & Wilkins, 2002.
- [2] A.I. González-González, M. Dawes, and J. Sánchez-Mateos et al. Information Needs and Information-Seeking Behavior of Primary Care Physicians. *Annals of Family Medicine* 5: 345-352, 2007.
- [3] Google Health homepage. <http://www.google.com/Top/Health>, 2007.
- [4] J. Groopman. How Doctors Think. Houghton Mifflin Company, 2007.
- [5] Healthline homepage. <http://www.healthline.com>, 2007.
- [6] W.R. Hersh, and D.H. Hickam. How Well do Physicians Use Electronic Information Retrieval Systems? A Framework for Investigation and Systematic Review. *JAMA* 280: 1347-1352, 1998.
- [7] G. Luo. iMed: An Intelligent Medical Web Search Engine. Available at <http://pages.cs.wisc.edu/~gangluo/imed.pdf>, 2008.
- [8] G. Luo, C. Tang. On Iterative Intelligent Medical Search. SIGIR 2008.
- [9] G. Luo. Intelligent Output Interface for Intelligent Medical Search Engine. AAAI 2008: 1201-1206.
- [10] G. Luo, C. Tang, and H. Yang et al. MedSearch: a Specialized Search Engine for Medical Information Retrieval. CIKM 2008.
- [11] Medstory homepage. <http://www.medstory.com>, 2007.
- [12] C. Sherman. Curing Medical Information Disorder. <http://searchenginewatch.com/showPage.html?page=3556491>, 2005.
- [13] SearchMedica - The GPs search engine. <http://www.searchmedica.co.uk/searchmedica/EUIHomeAction.do>, 2007.
- [14] J. Williams. When Expert Systems are Wrong. ACM SIGBDP Conference on Trends and Directions in Expert Systems 1990: 661-669.
- [15] D.L. Kasper, E. Braunwald, and A. Fauci et al. Harrison's Principles of Internal Medicine, 16th Edition. McGraw-Hill Professional, 2004.
- [16] P.M. Healey, E.J. Jacobson. Common Medical Diagnoses: An Algorithmic Approach, 2nd Edition. W.B. Saunders, 1994.
- [17] R.H. Sellar. Differential Diagnosis of Common Complaints, 4th Edition. W.B. Saunders, 2000.
- [18] American Medical Association Family Medical Guide, 4th Edition. John Wiley & Sons, 2004.
- [19] A.L. Komaroff. Harvard Medical School Family Health Guide. Free Press, 2004.