

# Intelligent Personal Health Record: Experience and Open Issues

Gang Luo · Chunqiang Tang · Selena B. Thomas

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

{luog, ctang, selenat}@us.ibm.com

**Abstract** Web-based personal health records (PHRs) are under massive deployment. To improve PHR's capability and usability, we previously proposed the concept of intelligent PHR (iPHR). By introducing and extending expert system technology and Web search technology into the PHR domain, iPHR can automatically provide users with personalized healthcare information to facilitate their daily activities of living. Our iPHR system currently provides three functions: guided search for disease information, recommendation of home nursing activities, and recommendation of home medical products. This paper discusses our experience with iPHR as well as the open issues, including both enhancements to the existing functions and potential new functions. We outline some preliminary solutions, whereas a main purpose of this paper is to stimulate future research work in the area of consumer health informatics.

**Keywords** Search engine · Personal health record · Expert system · Home medical product · Nursing activity


## 1. Introduction

Web-based personal health records (PHRs) are being widely deployed by a few major Internet companies, including Google [42], Microsoft [43], and Revolution Health [48]. Using PHRs, ordinary consumers can actively manage their medical records and ultimately their health through a Web interface. However, existing PHRs have limited intelligence and can meet only a small subset of users' healthcare needs. To improve PHR's capability and usability, we previously proposed the concept of an intelligent PHR (iPHR) [40, 41] by introducing and extending expert system technology and Web search technology into the PHR domain.

iPHR is intended to be a centralized portal that automatically provides users with comprehensive and personalized healthcare information to facilitate their daily activities of living. Since consumers typically lack medical knowledge and cannot identify appropriate keywords to search healthcare information, iPHR extensively uses medical knowledge to (1) anticipate users' needs, (2) guide

users to provide the most important information about their medical condition, and (3) automatically form queries.

Fig. 1 shows iPHR's architecture. Using medical knowledge, the expert system converts information in the PHR into a set of "search guide information" that reflects the user's medical condition and healthcare needs. This search guide information serves as seeds for the Web search engine to retrieve personalized healthcare information. iPHR also includes our intelligent medical Web search engine iMed [32, 33, 34, 37, 38], which uses disease diagnosis knowledge to help the user find disease information related to her medical condition.



**Fig. 1** Architecture of iPHR.

At present, our iPHR system provides three functions: guided search for disease information [32, 33, 34, 37, 38], recommendation of home nursing activities (HNAs) [39], and recommendation of home medical products (HMPs) [35, 36, 41]. These functions currently cover about one thousand health issues in various categories. During development and operation of iPHR, we obtained valuable experience and identified some open issues, including both enhancements to the existing functions and potential new functions. As iPHR is a relatively new field, this paper summarizes our experience, discusses the open issues, and outlines some preliminary solutions. One main purpose of this paper is to stimulate future research work in the area of consumer health informatics [16].

This paper is an extended version of our previous IHI'10 conference paper [77], and includes new material on automatically identifying contraindicated HNAs in Section 4.1, on figuring out all of the relevant nursing interventions and HNAs for each health issue in Section 4.2 (Problem 4 and the subsection entitled "Proposed solution for

compiling relevant HNAs”), and on automatically identifying contraindicated HMPs in Section 5.1.

The rest of the paper is organized as follows. Section 2 gives an overview of iPHR. Sections 3-6 discuss the open issues in iPHR. Section 7 presents related work. Section 8 concludes this paper.


## 2. Overview of iPHR

In this section, we provide an overview of iPHR. One major challenge in consumer health informatics [16] is that consumers are often uninformed of the medical problems they face and the related medical terminology (e.g., panophthalmitis). As a result, a key design principle of iPHR is to use medical knowledge to anticipate users’ healthcare needs and automatically form queries whenever possible. The details of our design rationale and implementation techniques are described in our previous publications [32, 33, 34, 35, 36, 37, 38, 39, 40, 41].

iPHR provides its intelligent functions to users through a Web interface. There are a few buttons on the right side of the main Web page of iPHR. Each button corresponds to a different, intelligent function. If the user clicks a button, she is directed to a page specifically designed for the corresponding function. Our iPHR system currently provides three functions: guided search for disease information, recommend HNAs, and recommend HMPs. The rest of this section summarizes these three functions.

### 2.1 Guided search for disease information

iPHR includes our intelligent medical Web search engine called iMed [32, 33, 34, 37, 38], which introduced and extended expert system technology into the search engine domain. iMed uses medical knowledge and an interactive questionnaire to guide users to provide the most important information about their medical condition and subsequently form queries.



**Fig. 2** The diagnostic decision tree for the symptom “chest pain” [11].

iMed leverages its built-in medical knowledge in the form of diagnostic decision trees written by medical professionals [11]. As shown in Fig. 2, each diagnostic decision tree corresponds to either an objective sign (e.g., hypertension) or a subjective symptom (e.g., fatigue). 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 is a topic (typically a disease) potentially relevant to the user’s medical condition.

iMed uses diagnostic decision trees to help the user form queries. The user first selects one or more signs and symptoms from a list of known signs and symptoms. Then iMed asks questions related to these selected signs and symptoms. Based on the user’s answers to the questions, iMed navigates the corresponding diagnostic decision trees and reaches multiple topics that are potentially relevant to the user’s medical condition. iMed displays both the search results and the most common aspects of these topics as a hierarchy, whose first level and second level are for topics and aspects, respectively. For each of these topics, iMed automatically uses the topic name as a query to retrieve some related Web pages. iMed also presents a predetermined set of aspects (e.g., symptom, diagnosis, treatment, and risk factor). If the user clicks a particular aspect of the topic, iMed automatically combines the aspect name and the topic name into a query and uses that query to retrieve some Web pages related to this aspect of the topic.

For example, Fig. 2 shows the diagnostic decision tree in Collins [11] for the symptom “chest pain.” If “chest pain” is the only symptom chosen by the user, iMed’s first question is “Is the pain constant or intermittent?” If the user answers “intermittent” to this question, iMed’s next question is “Is the pain increased by breathing?” If the user answers “yes” to the second question, iMed reaches multiple topics including pleurisy.

Since most users have little medical knowledge, they frequently encounter challenges and need assistance during the entire disease information search process. iMed provides various kinds of suggestions to facilitate the search process. First, iMed suggests diversified and related medical phrases to help the user quickly digest search results and refine his inputs. Second, iMed suggests signs and symptoms related to the user’s medical condition. Third, iMed suggests alternative answers to its questions when the user suspects that he answers questions incorrectly.

### 2.2 Recommending home nursing activities

Based on the user’s health issues, iPHR can automatically recommend HNAs (home nursing activities) [39]. Each nontrivial HNA is made clickable for the user to find various, detailed implementation procedures for it on the Web. This HNA recommendation function is made possible by using standardized nursing languages [30].

Over two decades and the efforts of thousands of nurses, the nursing informatics community has systematically organized nursing knowledge into several standardized nursing languages [30]. Among them, iPHR’s knowledge base has incorporated NANDA-I nursing diagnoses and NIC nursing interventions, which cover the full range of the nursing domain. NANDA-I and NIC are the acronyms for

North American Nursing Diagnosis Association International and Nursing Interventions Classification, respectively.

Each health issue links to one or more *NANDA-I nursing diagnoses*, which are clinical judgments about individual, family, or community responses to actual or potential health problems [1]. Every nursing diagnosis typically links to a list of 10 or more *NIC nursing interventions* [6, 30], which are treatments that can be performed to enhance patient/client outcomes. Each nursing intervention includes a list of usually 10 to 30 *nursing activities* that are used to implement it. Using nursing diagnoses and nursing interventions as intermediate steps, we can link each health issue to multiple nursing activities, as shown in Fig. 3. These nursing activities represent the actions that nurses, patients, and caregivers can take to achieve desirable outcomes for this health issue. Because of its consumer-centric view, iPHR focuses on HNAs that patients and caregivers can perform at home or in the community.




Fig. 3 Linking health issues to nursing activities.

iPHR's HNA recommendation function works in the following way. iPHR automatically extracts from PHR the health issues related to the user, e.g., his current disease. For these health issues, iPHR uses the above method to find all linked HNAs and displays them as a prioritized hierarchy. For each aspect of a nontrivial HNA, iPHR's knowledge base stores a pre-compiled phrase as its HNA search guide information. A link is added into the displayed Web page for this aspect. If the user clicks this link, iPHR submits that phrase as a query to a large-scale medical Web search engine and returns the search results to him. This helps the user find various, detailed implementation procedures for this aspect.

For example, the health issue asthma links to the nursing diagnosis of *ineffective airway clearance*, which links to the nursing intervention of *asthma management* that includes the HNA "Coach in breathing/relaxation techniques." For the breathing aspect of this HNA, the top results retrieved by the corresponding phrase "asthma breathing techniques" include: (1) the Buteyko method for breathing ([www.correctbreathing.com](http://www.correctbreathing.com)), (2) two new breathing exercises for asthma ([www.sciencedaily.com/releases/2008/05/080528095853.htm](http://www.sciencedaily.com/releases/2008/05/080528095853.htm)), and (3) the book "Free Your Breath, Free Your Life" teaching breathing techniques ([www.authenticbreathing.com/asthma.htm](http://www.authenticbreathing.com/asthma.htm)).

### 2.3 Recommending home medical products

Based on the user's health issues, iPHR can automatically recommend HMPs (home medical products) [35, 36, 41] by leveraging its built-in treatment knowledge and nursing knowledge. More specifically, for each HNA, a nurse pre-

compiles a set  $S_A$  of phrases as its HMP search guide information and stores  $S_A$  in iPHR's knowledge base. Each phrase in  $S_A$  provides one way of retrieving HMPs related to this HNA. For each health issue (e.g., disease, symptom, surgery), iPHR stores a list of search guide phrases pre-compiled using disease/symptom treatment knowledge. These phrases serve to bridge the semantic gap between the literal meaning and the underlying medical meaning of the health issue.

- < Medical Supplies & Equipment
- < **Daily Living Aids**
- \* Bath & Body Aids (7159)
- \* Eating & Drinking Aids (1438)
- \* Low Vision Aids (1235)
- \* Dressing Aids (1086)
- \* Low Strength Aids (671)
- \* Medication Aids (610)
- \* Hearing Aid Accessories (303)
- \* Ramps (286)
- \* Hearing Aids (226)
- \* Telephones (87)
- \* Others (856)

Fig. 4 A sample navigation hierarchy constructed for the health issue muscular dystrophy.

iPHR automatically extracts from PHR the health issues related to the user, e.g., his current disease. For these health issues, iPHR uses the method described in Section 2.2 to find all relevant HNAs. The complete set  $G$  of search guide information includes both the HMP search guide information for those HNAs and the search guide phrases pre-compiled for these health issues using treatment knowledge. iPHR submits each search guide phrase in  $G$  as a query to a vertical search engine and retrieves some relevant HMPs. Then iPHR combines all retrieved HMPs together and returns them to the user using a navigation interface. On each search result Web page, a navigation hierarchy based on product categories is displayed on the left side (see Fig. 4) whereas recommended HMPs are displayed sequentially on the right side.

For example, the health issue muscular dystrophy links to the nursing diagnosis of *risk for falls*, which links to the nursing intervention of *fall prevention* that includes the following HNAs:

- (1) Provide sturdy, nonslip *step stools* to facilitate easy reaches;
- (2) Provide *nightlight* at bedside; and
- (3) Provide visible *handrails* and *grab bars*.

The HMPs of step stool, nightlight, handrail, and grab bar are all relevant to muscular dystrophy. For both muscular dystrophy and its symptoms, neither their names nor their treatment methods may appear in the Web pages describing these HMPs. However, using nursing knowledge, iPHR can recommend these HMPs to muscular dystrophy patients.

The three existing functions in our iPHR system provide valuable healthcare information to facilitate users' daily activities of living. Nevertheless, many issues remain open in iPHR, in terms of enhancing the existing functions and providing new functions. In the following sections, we discuss these open issues and outline some preliminary solutions.

Many of these open issues are identified from feedback of iPHR users as well as from our own experience of developing and operating iPHR. In addition, the following activities have contributed to our thinking. First, communicating with patients, caregivers, and healthcare professionals. Second, observing people at home and in various medical settings, such as hospitals, rehabilitation centers, and nursing homes. Third, reviewing online medical discussion forums. Fourth, reading medical books and articles for the general public, introduction chapters in medical and nursing textbooks, and published patient diaries. We believe that these open issues provide excellent opportunities for future research work.

### 3. Open issues in guided search for disease information

In this section, we discuss the open issues in iPHR's function of guided search for disease information.

#### 3.1 Likelihood of having a disease

Online health information is valuable to consumers. However, it can also cause cyberchondria, a common phenomenon reducing people's quality of life [56]. After viewing online medical search results, a layman tends to panic unnecessarily and worry about having a serious but rare disease that is unlikely to apply to him. To prevent cyberchondria, iMed can present the incidence rate of each disease displayed in its search results. The incidence rate of a disease is the number of its new cases per 1,000 people per year and reflects the probability of developing it [12, 13]. To help the user better understand this number, we can draw an analogy, such as this number is roughly equivalent to the likelihood of encountering a shark attack.

Although this method helps, it is still insufficient for preventing cyberchondria. The user tends to think that he has the symptoms of a rare disease. As a result, he believes that his likelihood of having this disease is much higher than its incidence rate among the general population. In fact, inexperienced medical students also tend to make similar mistakes, partly due to human nature's inclination to focus on extreme cases [24, page 59].

A plausible solution to this problem is to show the conditional probability of having this disease given one's exact symptoms. Nevertheless, unlike precise test results, some symptoms can be fuzzy. A layman frequently misjudges his objective symptoms [8, 37] due to a lack of medical training. Without knowing the exact symptoms, we cannot compute this conditional probability.

Another solution to this problem that uses more information is to present the conditional probability of having this disease given the subjective information that one provides in iMed's questionnaire, including both the symptoms that he thinks he has and his answers to iMed's questions. iPHR can provide training data to estimate this conditional probability. However, it remains to be seen whether this conditional probability is close enough to his actual likelihood of having this disease and hence can have any practical value.

When a user invokes iMed to search disease information, his exact disease is unknown whereas iMed automatically records his subjective information. Usually within a few days, he will visit his doctor, get to know his exact disease, and record it in his PHR. The combination of his subjective information and his exact disease becomes one training data point. By keeping track of many users over some amount of time, iPHR will be able to accumulate a large training data set.

When we implement the above idea in our iPHR system, we encounter a few challenges. The user can have multiple symptoms and answer quite a few related questions posed by iMed. In other words, his subjective information usually contains many parts with strong dependency among themselves. Consequently, estimating this conditional probability is a high-dimensional sparse data problem. Given the limited amount of training data, we need to break this problem into multiple low-dimensional ones and resort to some smoothing techniques.

By default, we match the subjective information that the user provides in iMed's questionnaire with the disease that will be recorded in the next entry of his PHR. However, this match could be incorrect for two reasons. First, the user may search disease information on behalf of another person rather than himself. Second, the disease that will be recorded in the next entry of his PHR may not be what he is currently searching for. For example, the user's current symptoms disappear after some rest and his next visit to his doctor is due to other symptoms. To reduce the likelihood of obtaining an incorrect match, one possible method is to check the symptoms that he selects in iMed's questionnaire and the symptoms that will be recorded in the next entry of his PHR. If these two sets of symptoms are roughly consistent (not necessarily identical) and there is only a short delay before the next entry is recorded in his PHR, then our match is likely to be correct. Nevertheless, it is non-trivial to determine such consistency.

Besides showing the conditional probability of having a disease, it is also beneficial for iMed to present the conditional probability that the user has no disease and needs only some rest. If the symptoms of the user disappear after some rest, he may choose to not visit his doctor and hence no entry is recorded in his PHR for this event. Consequently, no training data is directly available for estimating the conditional probability of having no disease. It is possible that the above-mentioned method can be adapted to address

the problem. The status of having no disease this time could be implied by (1) a long delay before the next entry is recorded in his PHR or (2) strong inconsistency between the symptoms that he selects in iMed's questionnaire and the symptoms that will be recorded in the next entry of his PHR.

Due to the increasing lack of new doctors and the retirement of baby-boomer doctors [97], the interaction time between doctors and patients keeps shrinking. As this trend is expected to continue, it is important for patients to get prepared beforehand so that the limited amount of interaction time can be used effectively. Similar to what is mentioned in [96], iMed is intended to help people better prepare for doctors' appointments rather than perform self-diagnosis. In general, there is a large amount of information on the Web for any single disease. By presenting the conditional probabilities mentioned above, iMed can also help users focus on relevant diseases and do the preparation work more efficiently.

### 3.2 Aspect hierarchy

Our experience with iMed shows that it is helpful to display various aspects (e.g., symptom, diagnosis, treatment, and risk factor) for every disease potentially relevant to a user's medical condition. Currently, iMed displays a predetermined set of aspects that is the same for all diseases. Ideally, the set of aspects could vary from one disease to another based on the disease's unique characteristics. For example, diabetes could have the following aspects that do not apply to many other diseases:

- (1) Nutrition: Diabetic patients need to take special control of their diet.
- (2) Children: As shown in [23], diabetic children need special care due to their young age. For instance, it is a bad idea to reserve candy and give it to a diabetic child only when she is hypoglycemic (i.e., when her blood glucose level is lower than normal). She may be encouraged to make herself hypoglycemic in order to get candy.
- (3) Seniors: As shown in [14], diabetic seniors need special care due to their old age.

In general, a disease can have many aspects, e.g., based on factors such as age, pregnancy, gender, and disease stage. Each aspect can have its own sub-aspects. A sub-aspect can further have sub-sub-aspects, and so forth. The aspect hierarchy can be regarded as the table of contents of an encyclopedia on the disease. At each node of the aspect hierarchy, one or more targeted queries need to be formed to retrieve related Web pages. Since a semantic gap may exist between the name and the underlying medical meaning of a (sub-)aspect [41], the combination of the (sub-)aspect name and the disease name may not always form an appropriate query. In case of more than one query, these queries' search results need to be merged into a single, properly-ranked list [41].

The aspect hierarchy needs to be comprehensive. If a particular (sub-)aspect does not appear in it, the user may never be able to realize that there exists specific information about this (sub-)aspect of the disease. For example, the top Web pages retrieved using only the disease name may contain only general discussion about the disease but no such specific information.

So far we have compiled aspect hierarchies for a few diseases. We find it challenging to make the aspect hierarchy comprehensive and well-organized, to compile appropriate queries for the (sub-)aspects, and to properly merge multiple queries' search results for the same (sub-)aspect into a single list. To compile comprehensive (sub-)aspects as well as appropriate queries for these (sub-)aspects, a wiki tool [58] would be useful for enabling many volunteers to participate in the compilation process: medical professionals provide pointers to relevant medical text, medical librarians compile (sub-)aspects and corresponding queries, and consumers offer feedback. A Wiki tool could also enable many volunteers to participate in compiling search guide information for recommending HNAs as well as for recommending HMPs. To reduce the likelihood of accidentally omitting certain relevant aspects for some diseases, we could maintain a long reference list of all possible aspects from all diseases and check this list for every disease. To organize the aspect hierarchy well, we could resort to some statistical methods similar to the ones used in [6], e.g., similarity ratings and hierarchical clustering, for determining a good organization of nursing interventions. A comprehensive aspect hierarchy is often so large that users can have difficulty quickly finding their needed (sub-)aspects. To address this problem, we could develop a method for automatically suggesting the (sub-)aspects likely to be of interest to the user based on the information in her PHR. Some search tools could also be provided to help navigate the aspect hierarchy.

### 3.3 Ranking Web pages

Like any Web search engine, the Web page ranking method of an intelligent medical Web search engine critically affects its search result quality. We envision that the Web page ranking method could be improved in multiple ways, most of which could also apply to recommending HNAs and recommending HMPs.

First, as health knowledge continues to improve, some old health knowledge may no longer be valid or optimal. All other things being equal, ideally iMed should rank higher those health Web pages presenting up-to-date health knowledge. That is, like real-time search engines [50], iMed should consider both time [31] and traditional ranking criteria, such as Web site authority and similarity between query and Web page text, in its Web page ranking method. The most straightforward way to judge the content staleness of a health Web page is to use the timestamp of its HTML file if that file is static rather than dynamically generated.

However, this is neither the only way nor always the best way. On many health Web sites, e.g., WebMD [55], each Web page includes the latest date on which some healthcare professionals judged its content to be up-to-date. This date often reflects its content staleness more accurately than the timestamp of its HTML file.

Second, many health Web pages include images or multimedia content for illustration purpose, e.g., showing how to use a HMP, what a symptom looks like, and how to perform an exercise, a treatment procedure, or a nursing procedure. In general, multimedia content provides more information and is more effective at illustration than images. Without using any image or multimedia content, certain healthcare information can be rather difficult to understand with text-only description. One such example is how a hemiplegic patient, one side of whose body is paralyzed, can independently dress himself using only his unaffected hand [49]. For the same reason, modern health textbooks frequently include many pictures as well as attach multimedia CDs or DVDs.

All other things being equal, for a specific query, iMed could rank higher those health Web pages that include images or multimedia content illustrating related healthcare information. iMed could also rank higher Web pages including relevant multimedia content than Web pages including only relevant images. However, it is non-trivial to determine whether the multimedia content or the images on a Web page are relevant to this query, as their (alternative) text descriptions are often rather vague. One possible solution is to first use health knowledge to expand this query to include related health concepts, and then match the expanded query with both the (alternative) text descriptions of and the texts around the images or the multimedia content on the Web page.

### **3.4 Providing illustrative images and multimedia content**

At present, a considerable amount of healthcare information that is difficult to understand has no illustrative images or multimedia content available on the Web. For example, many HMPs have no corresponding videos to demonstrate how to use them. With only text descriptions, it is difficult for ordinary consumers to figure out how to use some of these HMPs. Ideally, for each topic or each (sub-)aspect of the topic, iMed should provide some illustrative images or multimedia content when applicable.

A picture is worth a thousand words. In addition to demonstrating complex usage scenarios, illustrative images and multimedia content can also generate an emotional response in consumers emphasizing the importance of certain health actions [19]. Otherwise, some consumers may be reluctant to perform these actions seriously.

For example, before performing a knee replacement surgery, physicians frequently ask the patient to lose some weight. If he can see a video showing rehabilitation after the surgery, he may be inspired to lose weight.

As a second example, consider gastric bypass surgery that treats obesity. The best treatment for obesity is with regular exercises and diet management over a long period of time. Many obese people view gastric bypass surgery as a quick fix to avoid exercises and dieting. A video showing the painful recovery and complications after the surgery may encourage compliance with suggestions to exercise and diet. Compared to only presenting the mortality rate and the complication rate of the surgery, showing the video is usually more effective at persuading them to make the right treatment choice.

As a third example, consider smoking that can cause ischemia (restriction in blood supply) to the extremities, resulting in tissue death and subsequently amputation of the affected area such as fingers or toes. Showing relevant amputation images can motivate smokers to quit. Similarly, a video of coronary artery bypass surgery and limitations on daily activities of living after the surgery may encourage smokers to quit.

Due to the large number of possible healthcare topics and their (sub-)aspects, it is labor intensive to produce all of the corresponding illustrative images and multimedia content. To address this problem, we can resort to collective contributions from users. Like YouTube [60], for each topic or each (sub-)aspect of the topic, iMed could allow users to upload images or multimedia content illustrating related healthcare information. Similarly, for each aspect of a nontrivial HNA or each HMP that is not easy to use, iPHR could allow users to upload illustrative images or multimedia content. The upload privilege should not be limited to healthcare professionals only. For example, as can be seen from consumer reviews on many HMPs on Amazon [4], consumers frequently come up with innovative methods of using HMPs for the same or new purposes. It is also possible that computer-sophisticated users can provide fairly fancy multimedia content, such as 3D animations allowing dynamic interactions so that people can view the targets from different angles using some iPhone-style rotation method. For accuracy and quality assurance, all uploaded material should be reviewed and managed by a dedicated team of healthcare professionals before it becomes publicly available.

### **3.5 Search result presentation**

We envision several other ways to improve the search result presentation of guided search for disease information. First, not every disease has a treatment or needs to be treated [21]. For example, prostate cancer often develops rather slowly and has little impact on an elderly patient's daily life for many years. The patient is likely to die from another disease before prostate cancer ever causes any significant problem. Knowing that one has a disease usually makes him stressed. This stress reduces his quality of life and can be harmful, as his psychological state can significantly affect his health [21]. As a result, if a person has a disease that either

has no treatment or treatment is not recommended, he may be better off not knowing that he has the disease. It would be desirable for iMed to allow the user to choose his level of psychological endurance. iMed presents different sets of diseases at various levels. At the default level, iMed withholds no information. At a specific, sensitive level, it is non-trivial to determine which diseases are appropriate for display. The decision made for a specific disease depends on multiple factors and certain computational results based on the information about the user in his PHR, e.g., the likelihood that he will die from another disease before this disease ever causes any serious problem.

Second, for most diseases, there is massive amount of online information, some at the basic level and others at more advanced levels. Without knowing the basic information first, a layman can have difficulty understanding the information that is more advanced. However, ordinary people often do not know the proper sequence of reading. They can easily become frustrated when they encounter advanced materials before they have a chance to see the basic materials. It would be desirable to incorporate sophistication levels into the aspect hierarchy so that for each disease, iMed's search results are presented like a curriculum. The user can start her reading at the basic level. If she has more time and is interested, she can continue to read the information that is more advanced.

#### 4. Open issues in recommending home nursing activities

In this section, we discuss the open issues in automatically recommending HNAs (home nursing activities).

##### 4.1 Automatically identifying contraindicated HNAs

A person often has multiple health issues. For example, 21% of Americans have multiple chronic conditions [54]. A HNA may be generally suitable for one health issue, but become undesirable in the presence of another health issue. In medicine, contraindication is the general term describing this phenomenon. Ideally, iPHR should not recommend contraindicated HNAs.

A contraindication is a health issue that makes a particular intervention undesirable [65]. For example, cancer is a contraindication for massage because cancer can spread through the lymphatic system. As massage increases lymphatic circulation, it may also potentially spread the cancer. Some contraindications are absolute, meaning that under no circumstance should the intervention be performed. For example, deep vein thrombosis is a health issue with blood clot formation in a deep vein. It is an absolute contraindication for massage because massage leads to stimulation of blood vessels and enhanced circulation, both of which may cause a blood clot to detach and occlude a blood vessel. Other contraindications are relative, meaning that the client is at a higher risk of treatment complications, but the risk may be outweighed by

other considerations or mitigated by other protective measures. For example, valve heart disease is a relative contraindication for the intervention of aerobic exercises because stress tolerance is low in severe valve heart disease.

Ideally, iPHR should automatically identify all contraindicated HNAs and provide an option for users to view them. For those absolutely contraindicated HNAs, iPHR should not return them to the user by default. For those relatively contraindicated HNAs, iPHR should either explicitly mark them as relatively contraindicated or not return them to the user by default. In the former case, preferably iPHR should provide some rationale or comments, which come from the original sources of the contraindication information and offer some insights into the nature of the concerns [65].

Automatic contraindication identification is valuable for not only ordinary consumers but also nursing professionals. The scope of nursing is so extensive that each nurse typically knows only a small part of it and is familiar with some but not all HNAs. Moreover, even if a nurse is familiar with a particular HNA, she may not know every contraindication for it [65]. In fact, for similar reasons, some PHRs provide the function of automatically checking adverse drug interactions [62].

To the best of our knowledge, no existing healthcare software can automatically identify contraindicated nursing activities. This absence is partly due to availability of health knowledge in a usable format. It also can take quite a few years for a particular collection of systematically compiled health knowledge to be incorporated into healthcare software. The first systematic compilation of contraindication information did not appear until the year of 2006 [65]. This is different from the case of drug-drug interaction information, which the medical community had systematically compiled many years ago [79].

##### Proposed solution for identifying contraindicated HNAs

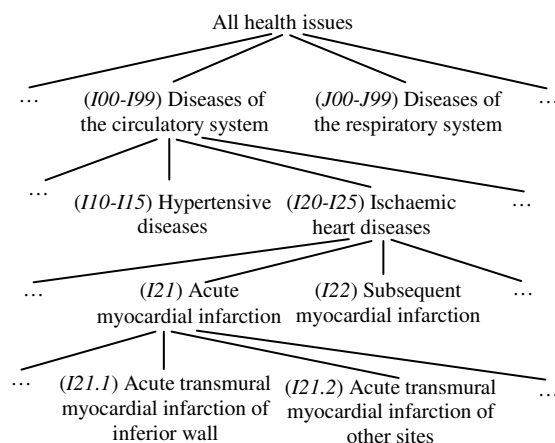


Fig. 5 The ICD-10 health issue hierarchy.

Our ongoing work in iPHR attempts to support automatic identification of contraindicated HNAs by leveraging the


medical ontology of International Classification of Diseases (ICD-10) [76, 95]. As shown in Fig. 5, the ICD-10 ontology organizes all health issues into a hierarchy. Each leaf node of the ICD-10 ontology corresponds to a “basic” health issue. Each non-leaf node of the ICD-10 ontology corresponds to a category of health issues.

We collect contraindication information from several sources (e.g., [65]). This information is stored in iPHR’s knowledge base in tuple form in the following way. In general, each piece of contraindication information shows that a health issue or a category of health issues  $H$  is a contraindication for a HNA  $A$ . We find the highest-level one or more ICD-10 nodes corresponding to  $H$ . These ICD-10 nodes are called the *contraindicated nodes* for  $A$ . They may or may not be leaf nodes. Each descendant node of any of them represents a more specific health issue that is also a contraindication for  $A$ . For example, if  $H$  is *ischaemic heart disease*, there is only one corresponding highest-level ICD-10 node: (I20-I25) ischaemic heart diseases. If  $H$  is *myocardial infarction*, there are two corresponding highest-level ICD-10 nodes: (I21) acute myocardial infarction and (I22) subsequent myocardial infarction. Each such contraindicated node  $N_o$  together with  $A$  form a tuple  $(A, N_o, R)$ , which is stored in iPHR’s knowledge base. Here,  $R$  contains the information of whether  $H$  is an absolute contraindication or a relative contraindication for  $A$  as well as some related rationale or comments. Essentially, all of these tuples form a relational database table. To facilitate search, an index is built on the HNA attribute of these tuples.

The main idea of our contraindication identification method is to check each linked HNA found with every current health issue of the user as well as every other health issue that the user cares about and has inputted. In theory, iPHR can identify all contraindicated HNAs if both its contraindication knowledge and its information about the user are complete. Our concrete approach is as follows.

As described in [41], iPHR automatically maintains a list  $L_{current}$  of current health issues of the user. Let  $L_c$  represent the list of health issues that the user cares about and needs iPHR to recommend related HNAs. Consider a particular health issue  $H \in L_c$ . The set  $S_H = L_{current} \cup L_c - \{H\}$  represents the other current health issues of the user and the other health issues that the user cares about and has inputted. As mentioned in Section 2.2, iPHR can use standardized nursing languages to find all HNAs linked to  $H$ . For each such HNA  $A$ , iPHR uses a **ICD-10 hierarchy-based method** to check whether any health issue in  $S_H$  is a contraindication for it. If so,  $A$  is a contraindicated HNA.

More specifically, for each health issue  $J \in S_H$ , we find the highest-level one or more ICD-10 nodes  $N_j$  corresponding to it. For any contraindicated ICD-10 node  $n$  for the HNA  $A$  stored in iPHR’s knowledge base, if any node in  $N_j$  is the same as  $n$ , an ancestor node of  $n$ , or a descendant node of  $n$  (see Fig. 6), then  $A$  is a contraindicated HNA because of the presence of  $J$ .



**Fig. 6** Identifying contraindicated HNAs using the ICD-10 hierarchy.

Let  $C_H$  represent the set of HNAs that are linked to the health issue  $H \in L_c$  but contraindicated for one or more health issues in the set  $S_H$ . The union of all such  $C_H$ ’s,  $\bigcup_{H \in L_c} C_H$ , represents the set of all contraindicated HNAs. If a HNA is absolutely contraindicated for one health issue and relatively contraindicated for another health issue, it should be regarded as absolutely contraindicated for the purpose of automatically identifying contraindicated HNAs.

When implementing automatic identification of contraindicated HNAs, the most difficult task we encountered is to obtain comprehensive information about contraindications. This information directly affects the quality of identification results but is labor intensive to compile manually for the following reason.

Contraindication information is currently scattered in numerous sources, such as medical journal articles, medical textbooks, and health Web sites, and frequently not indexed in the back of medical textbooks [65]. To the best of our knowledge, the only systematic compilation of contraindication information available so far is Batavia’s book [65], which covers about 100 interventions in a single area - physical rehabilitation. For a HNA not covered in this book, individual sources often contribute somewhat different contraindications for it (see part I of [65] for a detailed discussion). Consequently, to obtain a comprehensive set of contraindications for even a single HNA, we usually need to check many sources.

At present, we are compiling contraindications for the HNAs not covered in [65] using a method similar to that described in the Preface of [65]. First, we check medical textbooks. Second, we search both Web sources and medical journal databases. Besides the keywords of the HNA, our search also uses the following keywords: adverse effect, adverse event, caution, complication, contraindication, danger, harm, iatrogenic, precaution, risk, and safety. It would be interesting to investigate whether any semi-automatic method can help speed up this compilation process [35].

## 4.2 Gap between standardized nursing languages and HNA recommendation



Both NANDA-I nursing diagnoses and NIC nursing interventions were originally developed for purposes such as billing and standardization rather than recommending HNAs [1, 6, 30]. As a result, the properties of nursing diagnoses and nursing interventions only partially fulfill the requirements of recommending HNAs. This causes some problems, most of which also similarly exist for recommending HMPs.

#### **Problem 1: Lack of clarification on multi-linkage**

A nursing intervention can link to multiple nursing diagnoses. This design is used to ensure that for each health issue, all relevant nursing interventions can be found via linking through nursing diagnoses. Nevertheless, this design also implies that the same nursing intervention can link to a health issue multiple times, each time via a different nursing diagnosis. In some cases, such multi-linkage means that this nursing intervention should be given a higher priority. In other cases, such multi-linkage offers no extra information about this nursing intervention. However, the standardized nursing languages do not specify which case falls into which circumstance. At present, we assume that all cases belong to the first category of cases, i.e., multi-linkage means higher priority. When recommending HNAs, iPHR sorts all retrieved nursing interventions into a prioritized list under this assumption [39]. This solution is acceptable, but imperfect due to the assumption. A similar problem exists when linking HNAs to a nursing diagnosis through nursing interventions, as a HNA can appear in multiple nursing interventions.

#### **Problem 2: Lack of details in the description of HNAs**

Many HNAs are described at a high level without providing sufficient details. This makes it difficult to compile a comprehensive set of HNA search guide phrases. For example, a HMP that can help implement a HNA may not mention this usage in its description. Without putting the HMP name into the HNA search guide phrase, the top search results for that HNA may exclude HNA implementation procedures using this HMP. As a concrete example, consider the HNA “help the client with energy conservation.” One method to conserve energy is to use a shower chair when showering [75]. However, the HMP name *shower chair* is not mentioned in the description of this HNA. Without resorting to the keywords “shower chair,” none of the top search results for this HNA will be related to using a shower chair to help conserve energy. Due to the large amount of nursing knowledge scattered in many nursing resources, it would be labor intensive to find the necessary details for all HNAs.

#### **Problem 3: Context sensitivity**

Certain HNAs are meaningful only within the context of a concrete health issue. For them, we cannot directly compile HNA search guide phrases suitable for all possible health issues. Rather, we have to compile HNA search guide

phrases within the context of the concrete health issue. For example, the nursing intervention of *teaching: disease process* includes the HNA “Describe possible chronic complications, as appropriate.” For this HNA, we need to compile a separate HNA search guide phrase for each individual health issue: “chronic complications” plus the health issue name. In the particular case of asthma, the HNA search guide phrase would be “chronic complications asthma.”

Furthermore, some of these HNAs can become more specific if they are within the context of a concrete health issue. Otherwise, it can be difficult to figure out what could be done for them. For example, the health issue asthma links to the nursing diagnosis of *deficient knowledge*, which then links to the HNA “provide information to support self-management.” This HNA, however, does not provide any concrete information about self-management. In the context of asthma, this HNA should be rephrased more precisely as “instruct client in ways to maintain respiratory health: remain indoors as much as possible when air pollution levels are high; avoid extremes in hot and cold weather; avoid prolonged close contact with persons who have respiratory infection; ...” [75] It would be labor intensive to handle this case adequately because many health issues need to be manually handled for each such HNA.

#### **Problem 4: Irrelevant indirect linkage**

Nursing diagnoses are intentionally designed to be of a manageable number. This causes some irrelevant nursing interventions to link to a health issue. More specifically, a health issue  $H$  links to a set of nursing interventions through nursing diagnoses. This set includes not only all nursing interventions relevant to  $H$  but also some nursing interventions irrelevant to  $H$ . For instance, both  $H$  and another health issue  $H'$  link to a nursing diagnosis  $D$ , which links to two nursing interventions  $I$  and  $I'$ . It is possible that  $I$  is relevant to  $H$  but not  $H'$ , and  $I'$  is relevant to  $H'$  but not  $H$ . As a concrete example, the health issue conjunctivitis (pink eye) links to the nursing diagnosis of *acute pain*, which links to the nursing intervention of *patient-controlled analgesia (PCA) assistance* that is irrelevant to conjunctivitis. Similarly, some irrelevant HNAs can link to a nursing diagnosis because nursing interventions are intentionally designed to be of a manageable number.

When managing a patient with a specific health issue, nurses are trained to use their nursing knowledge to manually select only the relevant nursing interventions and HNAs [1, 6, 30]. However, most layman users of iPHR have no formal training in nursing and hence have difficulty differentiating the relevant nursing interventions and HNAs from the irrelevant ones. To enable iPHR to display only the relevant information, we need to figure out all of the relevant nursing interventions and HNAs for each health issue beforehand and record this information in iPHR's knowledge base. In other words, we need a comprehensive, standardized home nursing care plan for each health issue.

Documented nursing care for a health issue is based on a comprehensive, standardized nursing care plan that includes nursing interventions and nursing activities relevant to this health issue [70, 75, 78]. Frequently, each nursing activity is accompanied by some rationale showing why it is appropriate for this health issue and can produce desired response. Similar to evidence-based nursing [63, 69, 71], such rationales can help the user of iPHR figure out the nursing activities suitable for him and become motivated to perform them. Hence, it is desirable for iPHR to display next to each recommended HNA the corresponding rationale, if any.

The nursing community has compiled standardized nursing care plans for several hundred, but not all, health issues and published them as nursing textbooks [70, 75, 78, 84, 85, 86, 87, 88, 89, 90, 91]. Several nurses have also set up a few Web sites dedicated to compiling standardized nursing care plans for additional health issues [82, 83]. However, the existing standardized nursing care plans do not fully serve our purpose, as they typically are tailored to the hospital setting and include some, but not all, of the relevant HNAs.

To recommend HNAs with high accuracy, we need to figure out the relevant HNAs for each health issue. A nursing intervention that includes no relevant HNA is irrelevant to the health issue. At a first glance, this task looks daunting. There exist thousands of health issues whereas a health issue frequently links to hundreds of HNAs [1, 6, 30]. Even a single health issue can take a nurse quite some time to check all its linked HNAs and mark the irrelevant ones. Moreover, a nurse often has difficulty judging the relevance of a HNA to a health issue, as she knows only a small part of the extensive scope of nursing.

### ***Proposed solution for compiling relevant HNAs***

To reduce the amount of labor needed by this task, we adopt the following strategies. First, we reuse as much information that has already been compiled as possible. Second, we (semi-)automate as many parts of the relevance checking process as possible. Third, we use Web searches and medical journal database searches to help make up the knowledge deficit of the nurse who performs this task.

For each health issue, we figure out the relevant HNAs in the following way. Using the method described in Section 2.2, we find all of the HNAs linked to this health issue. For each such HNA, we check whether it is included in any of the existing standardized nursing care plans for this health issue, if any [70, 75, 78, 84, 85, 86, 87, 88, 89, 90, 91]. If so, it is relevant. Otherwise, we continue to check whether this health issue is a contraindication for it by using the ICD-10 hierarchy-based method described in Section 4.1. If it is contraindicated, it is irrelevant. Otherwise, we continue to probe its relationship with this health issue by searching the Web and medical journal databases. There are three possible cases:

- (1) If we find a high-quality health Web page or a medical journal article mentioning that it can be used for this health issue, then it is relevant. The associated rationale, if any, is recorded.
- (2) If we find a high-quality health Web page or a medical journal article mentioning that it should not be used for this health issue, then it is irrelevant. This new piece of contraindication information is added into iPHR's knowledge base. The associated rationale or comment, if any, is recorded. Since false and misleading health information exists on the Web, the quality of the health Web page is important and can be automatically judged by many existing tools [80, 81].
- (3) If no useful information is found via search, the nurse has to judge whether this HNA is relevant according to her own nursing knowledge and nursing practice experience. If she thinks this HNA is relevant, she tries to provide some rationale if possible. If she is uncertain, we stay on the conservative side and assume that this HNA is irrelevant by default.

To probe the relationship between a HNA and a health issue via search, one or more keyword queries are used. Each query includes both the health issue name and some keywords of the HNA. All of these queries are tried one by one until either some useful information is found or all of these queries have been exhausted.

For Web search, we require all keywords of the query to appear in the text of the Web page, e.g., by putting "allintext:" at the beginning of the query according to the syntax of the Google search engine [73, 74]. We observe that the anchor text of links to a Web page usually cannot help identify the relationship between the HNA and the health issue. However, Web search engines often use both anchor text and Web page text for query matching purpose [66]. If all keywords of the query appear in the combination of this anchor text and the Web page text, by default this Web page appears in the query's search results. For example, the Web page mentions the HNA as a side effect and something else relevant to the health issue as its main topic. The health issue name appears in the anchor text of some link to the Web page. In this case, most likely the nurse cannot identify the relationship between the HNA and the health issue by reading this Web page. Nevertheless, if all such Web pages are allowed to appear in the query's search results, these Web pages often dominate the top search results and make it difficult for the nurse to discover the Web pages that contain useful information.

So far, we have compiled the relevant HNAs for a limited number of health issues. Even after the compilation task is finished for each health issue, this relevancy information still needs to be updated from time to time, as nursing interventions and HNAs keep evolving [1, 6] and new nursing knowledge keeps coming up. For example, a new nursing research result can change the current perception of whether a particular HNA is relevant to a health issue. It would be interesting to investigate whether

any semi-automatic method can facilitate this maintenance task, which needs to be performed periodically and is rather time consuming [35]. For instance, Google Alerts can run a keyword query periodically and automatically notify the user when the query produces new search results [72]. To help discover new nursing knowledge related to this maintenance task, one approach is to set up a Google Alert for every keyword query that is used to probe the relationship between a HNA and a health issue.

### 4.3 Presentation of recommended HNAs

We envision several other ways to improve the presentation of recommended HNAs, most of which could also apply to recommending HMPs. First, when managing a patient with a particular health issue, it is possible that not all HNAs relevant to this health issue are suitable for her. In fact, nurses are trained to use their nursing knowledge to manually select HNAs based on the concrete situation of the patient, such as age, pregnancy, gender, current medications, and other concurrent health issues. By using information in the PHR of the patient, iPHR could further individualize the recommended HNAs by removing the HNAs that are generally relevant to her health issue but unsuitable for her.

Second, it is possible that multiple HNA search guide phrases are compiled for an aspect of a HNA, each phrase corresponding to a different way of implementing this aspect. Some ways, such as those widely-recognized home remedies using kitchen materials [46], can be cheaper or easier to implement or usually more effective than the others. In this case, the search results of all compiled phrases need to be merged into a single, properly-ranked list [41]. The search results of the phrases corresponding to the preferred ways should tend to rank higher, whereas search result diversification among all compiled phrases needs to be provided so that the top few search results do not all come from the same phrase [41].

Third, many HNAs have associated risks, warnings, and rationales. It would be desirable to compile HNA search guide phrases in such a way that they can retrieve not only HNA implementation procedures but also associated risks, warnings, and rationales. For example, some exercises, if performed improperly, could hurt a body part. As another example, two HNAs could conflict with one another and hence should not be performed together. If a user does not know the rationale for a HNA, he may either implement the HNA improperly or not give enough priority to performing the beneficial actions included in the HNA. For example, the preferred approach to losing weight is through regular exercises and dieting over a long period of time so that a small amount of weight is lost monthly. This method works in the long term because most of the weight loss is fat. If one uses a “magic” method such as laxatives to lose weight quickly, most of his weight loss will be water that can come back as fast as it is lost [15]. As a second example, giving too much insulin to a diabetic patient can make his blood glucose

level drop to dangerously low levels. As a third example, Thompson and Thielke [53] reported that a dementia patient suffered from low blood pressure because one of her family members checked her blood pressure frequently and gave her too much medication to lower it.

## 5. Open issues in recommending home medical products

In this section, we discuss the open issues in automatically recommending HMPs (home medical products).

### 5.1 Automatically identifying contraindicated HMPs

Section 4.1 describes a method of automatically identifying contraindicated HNAs. That method can also be applied to recommending HMPs, with the goal of avoiding recommending the HMPs related to the contraindicated HNAs.

One simple approach is to not use the HMP search guide phrases pre-compiled for the contraindicated HNAs. This approach, however, has limitations. A HMP often can be used for several purposes corresponding to multiple HNAs. In this case, its description may include two HMP search guide phrases: a first one pre-compiled for a desired HNA and a second one pre-compiled for a contraindicated HNA. Even without using the second HMP search guide phrase, the first HMP search guide phrase alone still retrieves this HMP, which can be undesirable for the user.

#### *Proposed solution for identifying contraindicated HMPs*

For safety purposes, we use a conservative approach and do not recommend any HMP whose description contains at least one HMP search guide phrase pre-compiled for a contraindicated HNA. More specifically, iPHR’s knowledge base includes a list of all possible HNAs. Let  $M_c$  represent the list of health issues that the user cares about and needs iPHR to recommend related HMPs. For each health issue  $H \in M_c$ , we use the ICD-10 hierarchy-based method described in Section 4.1 to identify all HNAs  $R_H$  for which  $H$  is a contraindication. These HNAs may or may not be relevant to any health issue in  $M_c$ . To reduce run-time overhead, such computation for each health issue can be performed beforehand with the result stored in iPHR. The union of all such  $R_H$ ’s,  $R_c = \bigcup_{H \in M_c} R_H$ , represents the set of all HNAs contraindicated for  $M_c$ . Recall that  $S_A$  represents the set of HMP search guide phrases pre-compiled for a HNA  $A$ . The set  $S_c = \bigcup_{A \in R_c} S_A$  represents all HMP search guide phrases pre-compiled for the HNAs in  $R_c$ . By treating each HMP search guide phrase in  $S_c$  as a sentence level Boolean conjunctive query, we can use the method described in [41] to find every HMP whose description contains at least one HMP search guide phrase in  $S_c$ . All of these HMPs are not recommended to the user.

## 5.2 Product categories

As described in Section 2.3, when presenting HMP search results, iPHR uses a navigation hierarchy based on product categories [36]. iPHR currently reuses the product categories of an HMP shopping Web site (Amazon [4]). These general product categories are designed for all health issues. They may not always be the best ones for an individual health issue. It would be desirable to automatically tune product categories for individual health issues, possibly using medical knowledge. Moreover, various HMP shopping Web sites use different product categories. It would be beneficial to (semi-)automatically match the product categories from different HMP shopping Web sites so that the HMPs from multiple HMP shopping Web sites can be mixed together in a single navigation hierarchy.

## 5.3 Linking to detailed HMP usage descriptions

iPHR recommends HMPs by displaying links to HMP Web pages from a few selected HMP shopping Web sites [41]. For many HMPs, their descriptions on these HMP Web pages are not detailed enough for consumers to understand how to use them. In contrast, a large number of Web pages on other non-shopping Web sites contain detailed usage descriptions of them. For each such HMP (more precisely, the HMP search guide phrase that is the HMP name), ideally iPHR should also present links to those Web pages that provide detailed descriptions on how to use it. This would require one or more query phrases to be pre-compiled for it and stored in iPHR's knowledge base. Recall that iPHR recommends the HMP for one or more relevant health issues. For each of some of the relevant health issues, the links to detailed HMP usage descriptions can target it by including its name in the query phrases. For a HMP that can be used to address a specific health issue, a query phrase including both the HMP name and the health issue name can retrieve Web pages showing how to use this HMP to help manage the health issue.

## 5.4 HMP search result presentation

We envision several other ways to improve the presentation of HMP search results. First, iPHR currently ranks HMPs based on their relevance to the user's health issues [41]. It would be beneficial to simultaneously consider multiple factors, such as popularity, price, and relevance to the user's health issues, in iPHR's HMP ranking criteria. Second, like HealthPricer [27], for the same HMP, iPHR could compare its prices from different sellers. Third, many activities have related health issues, e.g., exposure of the skin when fishing, gardening, or swimming. These activities are not covered in nursing diagnoses textbooks and hence not included in iPHR. It would be desirable to compile linked nursing diagnoses for these activities so that iPHR can recommend corresponding HMPs.

## 6. Potential new functions

In this section, we discuss some potential new functions of iPHR that can be helpful to consumers.

### 6.1 Automatically pushing healthcare information to users

iPHR currently follows the pull model of information distribution. When a user feels that he needs some healthcare information, he logs into iPHR and invokes a function that can provide this information. However, the pull model has some limitations. Due to a lack of medical knowledge, the user is not always able to realize that some beneficial healthcare information exists. By the time he gets to know this information, it may already be too late. For example, an Alzheimer's disease patient should set up a living will during the early stage of this disease [44]. Once he reaches the later stage of this disease and frequently becomes confused, he cannot validly sign any document for making legal or medical arrangements according to his wishes. As another example, many people are busy and live an unhealthy lifestyle. For instance, they frequently consume a certain kind of food with too much sodium that their hearts and vascular systems can be damaged (the food may or may not taste salty), but they are unaware of that fact. This is particularly problematic for people with a family history of hypertension.

To address this problem, we need to complement the pull model with the push model. iPHR should anticipate the needs of the user and automatically push personalized healthcare information to him. That information can either address his current medical issues or help him prevent future medical problems through education. iPHR can follow the push model of information distribution in two ways. First, iPHR can periodically email the user newsletters containing such information. Second, similar to the way Amazon displays potentially relevant products on its homepage [4], iPHR can display such information on its main Web page for the user to see during each login. In either case, iPHR needs a built-in recommender system [2] to automatically identify the healthcare information of potential interest to the user.

The user is interested in healthcare information relevant to herself. To locate this information, iPHR needs to use the information in her PHR, such as her current disease and her family history. The user can also serve as caregiver of other people, such as her kids, her spouse, and her mother living in another city, and hence is interested in healthcare information relevant to those people. To find this information, iPHR needs to use their information in their PHRs, which is readily available if they all use iPHR. To avoid overwhelming the user, iPHR needs her authorization about what kind of healthcare information could be pushed to her, in a way similar to how publish-subscribe systems work [47]. Before iPHR can start pushing the healthcare

information relevant to a user to another user, e.g., her caregiver, authorization from both users is needed. Since iPHR has access to many people's PHRs, theoretically it could even detect the outbreak of some health issue among a group of users with similar characteristics and push corresponding healthcare information to these users.

For a health issue that lasts some time, such as a chronic disease, a surgery, pregnancy, or raising a baby, the corresponding healthcare information pushed by iPHR could evolve as the health issue progresses. For example, the pushed healthcare information could go from basic knowledge to advanced knowledge, from a disease itself to related daily living issues, from before-surgery preparation to after-surgery recovery, or from caring for a newborn to caring for a 12-month-old baby. As a concrete example, before a patient has oral surgery, iPHR could let him know that his gums will remain numb for a few hours after the surgery. Consequently, during that period, he should avoid certain kinds of drink that can be either too hot or too cold and hence damage his mouth. By tracking the reading history of the user, iPHR could avoid pushing redundant healthcare information to him.

To push relevant healthcare information to users, we can proceed as follows. We pre-compile a list of possible topics and store them in iPHR's knowledge base. For each such topic, we pre-compile one or more search guide phrases and store them in iPHR's knowledge base. Periodically, iPHR automatically submits these search guide phrases as queries to a large-scale medical Web search engine to retrieve some relevant Web pages for the topic. Healthcare professionals manually select the best Web pages, possibly using a Wiki tool [58], and store them in iPHR's knowledge base. iPHR can then push relevant Web pages to the user based on the topics of interest to her.

## 6.2 Online social networking

Online social networking has fundamentally changed the way people obtain information. Both general-purpose social networking sites (e.g., Facebook.com and Twitter.com) and vertical social networking sites (e.g., LinkedIn.com for business professionals) are highly successful. In the healthcare domain, PatientsLikeMe [45] has started to pick up momentum by enabling patients to share symptom and treatment information to learn from real-world outcomes. iPHR could provide online social networking functions to help users obtain relevant healthcare information [17].

For example, iPHR could use the information in PHRs to match users according to certain criteria, such as having the same health issue and being located in the same city, and randomly pair up matched users for text or video chat [9]. Users can chat about their common health concerns and exchange their experience or relevant information, e.g., beneficial exercise programs offered in a local health club. Since the paired-up users usually do not know one another, it is likely they could obtain new information from each other.

iPHR could also pair up users to play online games with medical benefits, e.g., using the Wii gaming console [20, 57].

As a second example, like Twitter (twitter.com), iPHR could provide miniblog function. To facilitate users with the same health issue to connect with each other, all miniblogs are partitioned on discussed health issues. A user can choose to make some of her health issues public [45], possibly under a cyber name rather than her real name, and post miniblogs about them. For each health issue, iPHR can identify those users posting popular miniblogs about it. When a user records in her PHR that she has this health issue, iPHR can automatically recommend her to follow those users. Healthcare professionals could also register on iPHR and post miniblogs to connect with their colleagues and patients. For example, a physician can post an anonymized special case that he recently encountered or some new information on a health issue that he recently obtained from a medical conference. iPHR can automatically recommend users with a specific health issue to follow healthcare professionals posting popular miniblogs about the health issue.

Some users may use online social networking to forward their own agenda and/or perform detrimental activities. For instance, a commercial vendor may hire a person whose job is to distribute biased messages that dramatically exaggerate the effectiveness of its products. To help prevent such activities from happening, iPHR's online social network functions need to allow approved healthcare professionals to publicly report the credibility of a user, e.g., by giving a credibility score and/or leaving comments.

## 6.3 Automatic link generation

When reading a health Web page, a layman user may not realize the implications of some phrases on it due to a lack of medical knowledge. It would be desirable for iPHR to use medical knowledge to automatically generate links for the phrases that have certain implications [59]. If the user clicks such a link for a phrase, she will be directed to one or more Web pages presenting detailed implications of this phrase, preferably in the same context that this phrase is mentioned on the health Web page. For example, consider a health Web page mentioning that diabetic patients should try to replace part of their sugar consumption with sugar substitutes. According to medical knowledge, certain kinds of sugar substitutes can increase blood glucose level and hence are unsuitable for diabetic patients [51]. Using the pre-compiled query "sugar substitutes safe diabetics," an automatically generated link for the phrase "sugar substitutes" can point to multiple Web pages (e.g., [51]) discussing which sugar substitutes are safe for diabetics.

## 7. Related work

### *Consumer-centricity vs. physician-centricity*

The concept of intelligent electronic medical record (EMR) was proposed more than a decade ago. However, existing

intelligent EMRs are *physician-centric* and their intelligence is used to facilitate physicians' daily professional tasks, e.g., inputting and summarizing patient information [61]. In contrast, iPHR is *consumer-centric* and its intelligence is used to facilitate consumers' daily activities of living [40].

The medical informatics community traditionally focuses its work on addressing physicians' needs. Recently, it has begun to work on consumer health informatics [16] to address consumers' healthcare needs. Our iPHR work falls into the category of consumer health informatics.

### ***Electronic medical record***

Cimino has built an Infobutton Manager for EMRs [10]. For each medical concept appearing in the EMR, the Infobutton Manager provides a fixed set of questions that physicians ask most often and uses manually pre-constructed queries to retrieve answers to these questions from certain resources in real time. Each answer is retrieved using one query. The Infobutton Manager recommends neither HMPs nor HNAs. In contrast, for each health condition, our iPHR often uses hundreds of search guide phrases simultaneously to retrieve HMPs. As a result, iPHR can retrieve a much more comprehensive set of relevant HMPs than traditional keyword search, irrespective of whether a long query or a short query is used in traditional keyword search.

In general, tailored material tends to be regarded as more relevant and hence have larger impact [92]. Several existing personalized patient health information systems can dynamically generate tailored hypertext pages explaining the disease, its symptoms, and possible treatment options usually via using rules [92, 93, 94]. Some of these systems use the user's medical history stored in the EMR to make simple inferences on concepts she is likely to be unfamiliar with. Hyperlinks are added to only those concepts [92]. None of these systems recommends HMPs or HNAs.

Farfan et al. use ontology to facilitate keyword search in EMR [18], by incorporating the fact that many EMR standards are XML-based and have a hierarchical format. The method proposed in [18] does not apply to retrieving HMPs because HMP description has no hierarchical format.

### ***Symptom/disease search***

Healthline, a major medical Web search engine, added a symptom search function in Feb. 2007 [25, 26]. The user selects one or more symptoms from a given symptom list, and then Healthline returns a list of diseases having all of these symptoms. This function, however, has some limitations. First, many diseases have the same symptom. The correct way to narrow down the list of diseases is to consider additional information (e.g., gender, age, race, patient occupation) [11]. It is difficult and time-consuming for the user to distinguish these diseases himself by checking their detailed descriptions. For example, as mentioned in Collins [11], more than forty diseases can cause abdominal swelling. We can distinguish these diseases by using the

location and other properties of the swelling rather than using any other symptoms. Second, symptom search cannot express the absence of certain symptoms, which can significantly reduce the length of the list of possible diseases. Third, symptom search disallows the user from inputting other useful information for identifying the possible diseases, such as exam results, medication and food that the patient has taken, and existing diseases. Finally, a patient can have several symptoms at the same time due to the presence of multiple diseases. Symptom search can only find the diseases having all of these symptoms, whereas these diseases can be completely different from the patient's diseases. Even worse, the returned disease set for multiple symptoms is often empty. Our previous work on iMed [32, 33, 34, 37, 38] attempted to address these limitations.

When presenting the search results of a disease, Microsoft's Bing search engine [5] provides a small number of aspects of the disease. The set of aspects is the same for all diseases.

DXplain [64] is a diagnostic decision support system designed for physicians. It can provide diagnostic hypotheses by letting the user input information about the patient's condition via first answering a few questions and then entering several medical terms. This input process is often challenging for ordinary consumers because of the difficulty in choosing appropriate medical terms.

Besides diagnostic decision trees, there are other medical decision trees for treatment selection and diagnostic test selection, respectively [67, 68]. A person needs a lot of medical knowledge to understand and use those medical decision trees. Consequently, they are suitable for building healthcare software for medical professionals but unsuitable for building healthcare software for ordinary consumers.

### ***Product search***

Many product search engines have been launched [22]. They use no medical knowledge and cannot automatically recommend HMPs tailored to consumers' medical condition and healthcare needs.

The HMP shopping Web site AllegroMedical offers the function of searching HMPs by conditions [3]. This function is available for fewer than 50 conditions. For each condition, the list of matching HMPs is manually constructed and incomplete. AllegroMedical also offers the trial function of searching HMPs by body parts [29]. For each body part, the function essentially performs keyword search on HMPs using a list of manually pre-compiled keywords. However, the list is incomprehensive and covers only a small subset of the HMPs related to the body part.

Hidola [28] provides two functions: guided search for disease information and search HMPs. When presenting search results on disease information, Hidola does not provide any aspect of the topics that are potentially relevant to the user's medical condition.

### **Recommender system and general information retrieval**

Traditional recommender systems recommend items based on various factors such as users' prior ratings, previous purchases, and profiles [2]. Our iPHR can be regarded as a recommender system that goes beyond these factors and recommends personalized healthcare information using user profile, medical knowledge, and nursing knowledge.

Personalized search is a current trend of Web search engines [52]. Existing personalized search techniques adjust search results based on a user's search history and desktop content. Those techniques are useful for general search. However, since they leverage neither the user's medical history stored in the PHR nor medical knowledge essential to obtain high-quality queries, they are less effective at providing useful, personalized healthcare information, which is the focus of iPHR.

In distributed information retrieval and meta-search engines [7], search results from multiple sources for the same query are merged together. In contrast, in our case of automatic HMP recommendation, we need to merge together the HMP Web pages retrieved for different topics (by various query phrases).

### **Home care nurses**

Besides iPHR, home care nurses also provide healthcare information to facilitate people's daily activities of living. The scope of nursing is so extensive that each nurse typically knows only a small part of it. Nevertheless, a person often has multiple health issues and requires a wide range of healthcare information (e.g., 21% of Americans have multiple chronic conditions [54]). Moreover, healthcare information keeps updating rapidly and no nurse can always keep up with the latest ones. For example, as medical knowledge and technology continue to improve, each year many new HMPs enter the market. We would expect iPHR to complement home care nurses in providing healthcare information, because its knowledge base stores a comprehensive set of nursing knowledge compiled by thousands of nurses whereas its search system can discover the latest healthcare information from the Web. For a similar reason, iPHR can provide more complete information on HMPs and HNAs than any single consumer-oriented health information book (e.g., the series of books entitled "The Comfort of Home" [44]) or medical Web site (e.g., WebMD [55]).

## **8. Conclusions**

Intelligent personal health record is a new and rapidly moving field. This paper presents the valuable experience and the open issues we identified in developing and operating the intelligent personal health record system iPHR. iPHR extensively uses medical knowledge and Web search technology to provide consumers with personalized healthcare information to facilitate their daily activities of living. We outline preliminary solutions to some of the open

issues and expect that much research work is still needed to address the open issues to a satisfactory degree. Hence, one main purpose of this paper is to stimulate future research work in the area of consumer health informatics.

Among the open issues presented, our limited experience suggests the following three issues as the most important ones with the promise of obtaining feasible solutions in the next few years: (1) automatically identifying contraindicated HNAs (Section 4.1), (2) figuring out all of the relevant nursing interventions and HNAs for each health issue (Problem 4 in Section 4.2), and (3) automatically identifying contraindicated HMPs (Section 5.1). Some preliminary solutions to these three issues are outlined in this paper.

## **References**

1. Ackley, B.J., and Ladwig, G.B., *Nursing Diagnosis Handbook: An Evidence-based Guide to Planning Care*, 8th ed. Mosby, 2007.
2. Adomavicius, G., and Tuzhilin, A., Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *TKDE* 17(6): 734-749, 2005.
3. AllegroMedical homepage. <http://www.allegromedical.com>, 2010.
4. Amazon homepage. <http://www.amazon.com>, 2010.
5. Bing homepage. <http://www.bing.com>, 2010.
6. Bulechek, G.M., Butcher, H.K., and Dochterman, J.M., *Nursing Interventions Classification (NIC)*, 5th ed. Mosby, 2007.
7. Baeza-Yates, R.A., and Ribeiro-Neto, B.A., *Modern Information Retrieval*. ACM Press/Addison-Wesley, 1999.
8. Billings, J.A., and Stoeckle, J.D., *The Clinical Encounter: A Guide to the Medical Interview & Case Presentation*, 2nd ed. Mosby, 1999.
9. Chatroulette homepage. <http://chatroulette.com>, 2010.
10. Cimino, J.J., Use, usability, usefulness and impact of an Infobutton manager. *Proceedings of AMIA'06*, pp. 151-155, 2006.
11. Collins, R.D., *Algorithmic Diagnosis of Symptoms and Signs: Cost-Effective Approach*. Lippincott Williams & Wilkins, 2002.
12. Conditions by incidence. <http://www.wrongdiagnosis.com/lists/incid.htm>, 2010.
13. Data & Statistics, Centers for Disease Control and Prevention. <http://www.cdc.gov/DataStatistics>, 2010.
14. Dunning, T., *Nursing Care of Older People with Diabetes*. Wiley-Blackwell, 2005.
15. Duyff, R.L., *American Dietetic Association Complete Food and Nutrition Guide*, 3rd ed. Wiley, 2006.
16. Eysenbach, G., Consumer health informatics. *BMJ* 320(7251): 1713-1716, 2000.

17. Eysenbach, G., Medicine 2.0: social networking, collaboration, participation, apomediation, and openness. *J Med Internet Res* 10(3): e22, 2008.
18. Farfan, F., Hristidis, V., and Ranganathan, A., *et al.*, XOntoRank: Ontology-aware search of electronic medical records. *Proceedings of ICDE'09*, pp. 820-831, 2009.
19. Gagliano, M.E., A literature review on the efficacy of video in patient education. *J Med Educ.* 63(10): 785-792, 1988.
20. Garra, K., Medical benefits of Wii. <http://newhorizonllc.net/company-news/38-news/71-medical-benefits-of-wii>, 2010.
21. Goetz, T., *The Decision Tree: Taking Control of Your Health in the New Era of Personalized Medicine*. Rodale Books, 2010.
22. Google product search homepage. <http://www.google.com/products>, 2010.
23. Hanas, R., *Type 1 Diabetes in Children, Adolescents and Young Adults: How to Become an Expert on Your Own Diabetes*, 4th ed. Class Publishing, 2009.
24. Hanson, W., *The Edge of Medicine: The Technology That will Change Our Lives*. Palgrave Macmillan, 2008.
25. Healthline homepage. <http://www.healthline.com>, 2010.
26. New Healthline symptom search dramatically improves one of the most popular online health research activities. [http://www.healthline.com/corporate/news/healthline\\_announces\\_symptom\\_search.html](http://www.healthline.com/corporate/news/healthline_announces_symptom_search.html), 2007.
27. HealthPricer homepage. <http://www.healthpricer.com>, 2010.
28. Hidola homepage. <http://www.hidola.com>, 2010.
29. Hood, C., Unique holiday gifts - fitness equipment. <http://dailylivingaids.wordpress.com/tag/unique-gifts>, 2007.
30. Johnson, M., Bulechek, G.M., and Dochterman, J.M., *et al.*, *NANDA, NOC, and NIC Linkages: Nursing Diagnoses, Outcomes, and Interventions*, 2nd ed. Mosby, 2005.
31. Li, X., and Croft, W.B., Time-based language models. *Proceedings of CIKM'03*, pp. 469-475, 2003.
32. Luo, G., Design and evaluation of the iMed intelligent medical search engine. *Proceedings of ICDE'09*, pp. 1379-1390, 2009.
33. Luo, G., Intelligent output interface for intelligent medical search engine. *Proceedings of AAI'08*, pp. 1201-1206, 2008.
34. Luo, G., Lessons learned from building the iMed intelligent medical search engine. *Proceedings of EMBC'09*, pp. 5138-5142, 2009.
35. Luo, G., On search guide phrase compilation for recommending home medical products. *Proceedings of EMBC'10*, pp. 2167-2171, 2010.
36. Luo, G., Navigation interface for recommending home medical products. *JMS*, to appear.
37. Luo, G., and Tang, C., On iterative intelligent medical search. *Proceedings of SIGIR'08*, pp. 3-10, 2008.
38. Luo, G., and Tang, C., Challenging issues in iterative intelligent medical search. *Proceedings of ICPR'08*, pp. 1-4, 2008.
39. Luo, G., and Tang, C., Automatic home nursing activity recommendation. *Proceedings of AMIA'09*, pp. 401-405, 2009.
40. Luo, G., Thomas, S.B., and Tang, C., Intelligent consumer-centric electronic medical record. *Proceedings of MIE'09*, pp. 120-124, 2009.
41. Luo, G., Thomas, S.B., and Tang, C., Automatic home medical product recommendation. *JMS*, to appear.
42. Metz, C., Google eyes Cleveland medical records. [http://www.theregister.co.uk/2008/02/21/google\\_health\\_records\\_private\\_launch](http://www.theregister.co.uk/2008/02/21/google_health_records_private_launch), 2008.
43. Microsoft HealthVault homepage. <http://www.healthvault.com>, 2010.
44. Meyer, M.M., and Derr, P., *The Comfort of Home: a Complete Guide for Caregivers*, 3rd ed. CareTrust Publications LLC, 2007.
45. PatientsLikeMe homepage. <http://www.patientslikeme.com>, 2010.
46. Prevention magazine editors. *The Doctors Book of Home Remedies*. Bantam, 2003.
47. Ramasubramanian, V., Peterson, R., and Sirer, E.G., Corona: A high performance publish-subscribe system for the World Wide Web. *Proceedings of NSDI'06*, 2006.
48. Revolution Health homepage. <http://www.revolutionhealth.com>, 2010.
49. Radomski, M.V., and Trombly, C.A., *Occupational Therapy for Physical Dysfunction*, 6th ed. Lippincott Williams & Wilkins, 2007.
50. Sullivan, D., What is real time search? Definitions & players. <http://searchengineland.com/what-is-real-time-search-definitions-players-22172>, 2009.
51. Tatum, M., Which sugar substitutes are safe for diabetics? <http://www.wisegeek.com/which-sugar-substitutes-are-safe-for-diabetics.htm>, 2010.
52. Teevan, J., Dumais, S.T., and Liebling, D.J., To personalize or not to personalize: modeling queries with variation in user intent. *Proceedings of SIGIR'08*, pp. 163-170, 2008.
53. Thompson, H.J., and Thielke, S.M., How do health care providers perceive technologies for monitoring older adults? *Proceedings of EMBC'09*, pp. 4315-4318, 2009.
54. Vogeli, C., Shields, A.E., and Lee, T.A., *et al.*, Multiple chronic conditions: prevalence, health consequences, and implications for quality, care management, and costs. *JGIM* 22 (Suppl 3): 391-395, 2007.
55. WebMD homepage. <http://www.webmd.com>, 2010.
56. White, R.W., and Horvitz, E., Cyberchondria: Studies of the escalation of medical concerns in Web search. *ACM Trans. Inf. Syst.* 27(4), 2009.
57. Wii Fit homepage. <http://wiifit.com>, 2010.
58. Wikipedia homepage. <http://www.wikipedia.org>, 2010.



59. Wilkinson, R., and Smeaton, A.F., Automatic link generation. *ACM Comput. Surv.* 31(4), 1999.
60. YouTube homepage. <http://www.youtube.com>, 2010.
61. Yang, K., Evens, M.W., and Trace, D., Improving the text summaries in the intelligent medical record system. *Proceedings of MAICS'06*, 2006.
62. American Dietetic Association. *Electronic medical records and personal health records: A call for the creation and inclusion of a nutrition dataset*. [http://anhi.org/Learning/PDFs/Other/EMR-PHR\\_Nutrition\\_Dataset\\_Paper\(6-29-09\).pdf](http://anhi.org/Learning/PDFs/Other/EMR-PHR_Nutrition_Dataset_Paper(6-29-09).pdf), 2009.
63. Bakken, S., Currie, L.M., and Lee, N.J., *et al.*, Integrating evidence into clinical information systems for nursing decision support. *I. J. Medical Informatics* 77(6): 413-420, 2008.
64. Barnett, G.O., Cimino, J.J., and Hupp, J.A., *et al.*, DXplain. An evolving diagnostic decision-support system. *JAMA* 258(1): 67-74, 1987.
65. Batavia, M., *Contraindications in Physical Rehabilitation: Doing No Harm*. Saunders, 2006.
66. Brin, S., and Page, L., The anatomy of a large-scale hypertextual Web search engine. *Computer Networks* 30(1-7): 107-117, 1998.
67. Collins, R.D., *Algorithmic Approach to Treatment*. Lippincott Williams & Wilkins, 1997.
68. Collins, R.D., *Algorithmic Selection and Interpretation of Diagnostic Tests*. Williams & Wilkins, 1998.
69. Demner-Fushman, D., Seckman, C., and Fisher, C., *et al.*, A prototype system to support evidence-based practice. *Proceedings of AMIA'08*, pp. 151-155, 2008.
70. Doenges, M., Moorhouse, M., and Murr, A., *Nursing Care Plans: Guidelines for Individualizing Client Care across the Life Span*, 8th ed. F.A. Davis Company, 2009.
71. Elhadad, N., Kan, M., and Klavans, J.L., *et al.*, Customization in a unified framework for summarizing medical literature. *Artificial Intelligence in Medicine* 33(2): 179-198, 2005.
72. Google Alerts homepage. <http://www.google.com/alerts>, 2010.
73. Google Guide Quick Reference: Google Advanced Operators (Cheat Sheet). [http://www.googleguide.com/advanced\\_operators\\_reference.html](http://www.googleguide.com/advanced_operators_reference.html), 2010.
74. Google homepage. <http://www.google.com>, 2010.
75. Haugen, N., and Galura, S.J., *Ulrich & Canale's Nursing Care Planning Guides: Prioritization, Delegation, and Critical Thinking*, 7th ed. Saunders, 2010.
76. International Classification of Diseases (ICD-10) homepage. <http://www.who.int/classifications/icd/en/>, 2010.
77. Luo, G., Tang, C., and Thomas, S.B., Intelligent personal health record: experience and open issues. *Proceedings of IHI'10*, pp. 326-335, 2010.
78. Rodgers, S., *Delmar's Medical-Surgical Nursing Care Plans*. Delmar Cengage Learning, 2007.
79. Tatro, D.S., *2011 Drug Interaction Facts: The Authority on Drug Interactions*. Lippincott Williams & Wilkins, 2010.
80. Wang, Y., and Liu, Z., Automatic detecting indicators for quality of health information on the Web. *I. J. Medical Informatics* 76(8): 575-582, 2007.
81. Wilson, P., How to find the good and avoid the bad or ugly: a short guide to tools for rating quality of health information on the internet. *BMJ* 324(7337): 598-602, 2002.
82. <http://nursingcrib.com/nursing-care-plan/>, 2010.
83. <http://www.nursingcareplan.blogspot.com/>, 2010.
84. Comer, S.R., *Delmar's Geriatric Nursing Care Plans*, 3rd ed. Delmar Cengage Learning, 2004.
85. Comer, S.R., *Delmar's Critical Care Nursing Care Plans*, 2nd ed. Delmar Cengage Learning, 2004.
86. Luxner, K.L., *Delmar's Maternal-Infant Nursing Care Plans*, 2nd ed. Delmar Cengage Learning, 2004.
87. Luxner, K.L., *Delmar's Pediatric Nursing Care Plans*, 3rd ed. Delmar Cengage Learning, 2004.
88. Gulanick, M., and Myers, J.L., *Nursing Care Plans: Diagnoses, Interventions, and Outcomes*, 7th ed. Mosby, 2010.
89. Swearingen, P.L., *All-in-One Care Planning Resource: Medical-Surgical, Pediatric, Maternity, and Psychiatric Nursing Care Plans*, 2nd ed. Mosby, 2007.
90. Schultz, J.M., and Videbeck, S.L., *Lippincott's Manual of Psychiatric Nursing Care Plans*, 8th ed. Lippincott Williams & Wilkins, 2008.
91. Varcarolis, E.M., *Manual of Psychiatric Nursing Care Planning: Assessment Guides, Diagnoses, Psychopharmacology*, 4th ed. Saunders, 2010.
92. Cawsey, A.J., Jones, R.B., and Pearson, J., The evaluation of a personalised health information system for patients with cancer. *User modeling and user-adapted interaction* 10(1): 47-72, 2000.
93. Al-Busaidi, A., Gray, A., and Fiddian, N., Personalizing web information for patients: linking patient medical data with the web via a patient personal knowledge base. *Health Informatics J.* 12(1): 27-39, 2006.
94. Bental, D.S., Cawsey, A., and Jones, R., Patient information systems that tailor to the individual. *Patient Educ Couns.* 36(2): 171-80, 1999.
95. ICD-10 ontology. [https://dkm.fbk.eu/index.php/ICD-10\\_Ontology](https://dkm.fbk.eu/index.php/ICD-10_Ontology), 2011.
96. Luo, G., Tang, C., and Yang, H., *et al.*, MedSearch: a specialized search engine for medical information retrieval. *Proceedings of CIKM'08*, pp. 143-152, 2008.
97. A future doctor shortage? A growing lack of new doctors could be American's next health crisis. <http://www.cbsnews.com/stories/2006/06/17/eveningnews/main1726479.shtml>, 2006.