In Patel V, Rogers R and Haux R (Eds.) 10th World Congress on Medical Informatics (MedInfo'2001), September 2 – 5, London. IOS Press, Amsterdam.

Patient Empowerment via 'Pushed' Delivery of Personalised Healthcare Educational Content Over the Internet

Syed Sibte Raza Abidi, Chong Yong Han and Samina Raza Abidi

Health Informatics Research Group, School of Computer Sciences, Universiti Sains Malaysia, Penang, Malaysia

Abstract

We present an Internet-based Personalised Healthcare Information (PHI) dissemination system. Information personalisation is guided by the individual's current health profile as recorded in his/her EMR. A PHI package is composed by intelligently selecting and synthesizing various topic-specific documents, each corresponding to some health parameter noted in the EMR. To ensure medical consistency, constraint satisfaction techniques are employed during the information selection phase. The resultant PHI package—covering both long-term and immediate health-maintenance requirements—can be pro-actively pushed to the individual via email, thereby ensuring the timely availability of situation-specific health maintenance information. The featured work is in line with the Malaysian Multimedia Super Corridor Telemedicine initiative and can serve as a test-bed to evaluate the effectiveness of PHI, system design and operational considerations for larger-scale deployment.

Keywords:

Patient Empowerment; Personalised Healthcare Information; Constraint Satisfaction Techniques; Push Technology; Medical Ontology; XML

Introduction

Focussed and timely accessibility to quality healthcare information is central to patient empowerment initiatives [1] [2] [3]. Easier access to healthcare information—at least measured in terms of information volume—is seemingly addressed by the proliferation of healthcare information Web-sites over the Internet [4] [5]—the so-called *E-Healthcare Portals*. Notwithstanding the relevance of web-mediated healthcare information dissemination, the drawbacks noted are that a non-specialist individual has to (a) the information available does not specifically focus on the individual's current health problems and needs, rather it is too generic in nature; (b) meticulously sift through volumes of healthcare information, in an unguided manner, to search for relevant information; (c) make 'value' judgments about the relevance of the available healthcare information with respect his/her current health needs; and (d) pro-act to acquire healthcare information—the onus is on the individual to collect health maintenance information.

To address the above-mentioned limitations, and in turn to enhance the efficacy of patient empowerment initiatives, we advocate the generation and pro-active delivery of *Personalised Health Information* (PHI) for each individual [6] [7] [8] [9] [10]. We propose that for maximum impact, PHI content comprising health education to health promotion material should be 'dynamically' compiled based on an individual's current health profile, thereby specifically addressing the individual's current health needs. This calls for the dynamic personalisation of generic healthcare information by taking into account the individual's (i) chronic (long term) healthcare needs (ii) episodic (short term) healthcare needs, and (ii) healthcare objectives. Furthermore, the currency of the PHI should be ensured by the pro-active 'push based' delivery of the right information to the individual at the right time.

To meet the abovementioned proposal, we present the design and functional characteristics of a *Personalised Healthcare Information Delivery System* (PHIDS) [6]. From a functional viewpoint, PHIDS is an intelligent system, that exhibits the following features: (a) System-initiated, periodic generation of PHI at scheduled time intervals; (b) Dynamic PHI composition based on the individual's current *Health Profile* (HP) derived from information taken from his/her Electronic Medical Record (EMR); (c) Overall PHI content derived from multiple, heterogeneous *Topic-specific Documents* (TD) that are systematically integrated to yield a *PHI package*; (d) PHI package's consistency validation by intelligently ensuring that the content of the constituent TD do not contradict each other—i.e. the satisfaction of mutual (co-existence) constraints between multiple information items; (e) Use of a knowledge-base of constraint rules provided by medical experts; (f) Medical ontology based organisation of health information; (g) Health information encoded as Extensible Mark-up Language (XML) files; and (h) Pro-active 'push-based' delivery of PHI over the Internet.

The PHI generated by PHIDS comprises three main sections: (1) Information about long term-term clinical conditions and management regimes; (2) Information about short-term therapy and rehabilitation associated with nonacute illness episodes; and (3) General healthcare education with a wellness maintenance connotation.

Functional Overview of PHIDS

The overall functionality of PHIDS can be divided into three main activities:

1. *Generation of an Up-to-date Health Profile* based on information contained in individual-specific EMR and characteristic information acquired via system-initiated Web-based consultations sessions.
2. *Composition of a PHI Package* by systematically amalgamating multiple TD based on the individual's most current HP.
3. *Delivery of the PHI Package* using both pull (client-mediated) and push (system-mediated) methods.

We next discuss in detail certain key functional issues pertaining to the abovementioned functionality of PHIDS.

Task 1:Generation of an Up-to-date Health Profile

In essence, the generation of an up-to-date individual-specific HP demands the collection and summarization of relevant and recent information from the individual's EMR. If needed, the automatically generated HP can be verified by the individual via a system-initiated Web-based consultation session. Figure 1 depicts the overall process flow for HP generation.

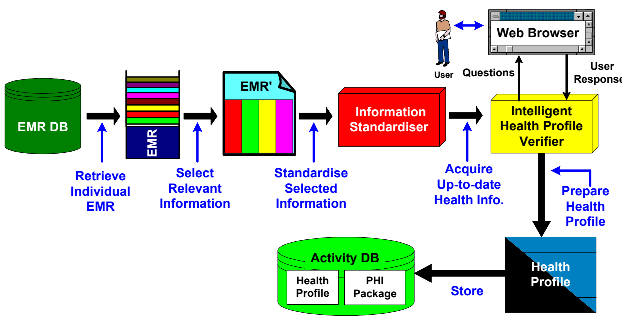

Figure 1 - The process flow for HP generation

Functionally, HP generation is carried out as follows:

(i) The temporal span of the HP is determined by making use of a *Health Profile Window (HPW)*—the retrospective timeframe spanning from the current date to some specified past date (typically set to 3 weeks) over which the EMR content would be examined for HP generation purposes.

(ii) Notwithstanding the information-rich content of an EMR, the HP is generated by selecting only a few relevant fields of the EMR—i.e. fields that refer to (a) acute diseases with corresponding therapeutic regime and (b) episodic encounters with healthcare providers.

(iii) The HP content is standardised with respect native terminological standards used to index healthcare information. The *Information Standardiser* module implements the Unified Medical Language Source (UMLS) meta-thesaurus and an ICD10 translator for both vocabulary and conceptual standardisation. The standardized HP content is considered as a *draft HP*.

(iv) The draft HP is presented to the individual via a WWW-based consultation session to be validated. The *Intelligent HP Verifier* module is designed to intelligently ask a series of relevant questions to validate the draft HP.

(v) Finally, the user-validated draft HP is deemed as the individual's 'most representative' HP and is used for the composition of the PHI package. The final HP is stored in the *Activity DB* for future references.

Task 2: Composition of the PHI Package

The dynamically composed PHI package—a comprehensive XML-based document—is derived by the systematic amalgamation of multiple TD, each with different levels of coverage and generality. Each TD is represented as a XML document. Our approach for generating PHI has two main tenets: (1) The collection of multiple candidate TD is based on a medical ontology that defines the hierarchical organisation of the *Healthcare Information Repository (HIR)*. (2) The selection of TD is determined by the 'intelligent' satisfaction of co-existence constraints between the candidate TD. Constraint satisfaction techniques are used to ensure that the aggregated information from multiple TD is medically consistent—i.e. upon aggregation the individual TD do not contradict each other, or lead to improper recommendations. We now briefly discuss the process workflow for PHI composition (as shown in Figure 2) with details about the modules used for PHI composition.

(i) Based on the individual's HP (generated earlier), a set of TD—where each TD corresponds to some medical concept within the HP—are collected from the HIR.

The HIR models a medical ontology derived from the hierarchical representation of concepts given in ICD10 coding and the index of UMLS. At the top level, information is categorised into the following categories: *Allergy, Diseases, Drugs, Lifestyle, Symptoms, Medical Dictionary* and *Medical Vocabulary*. These categories then further expand to multiple levels, from generic to specific health topics/concepts, as per the medical ontology. Each level of the ontology refers to a unique medical concept (or topic), where each concept is associated with a unique TD (an XML file) that contains details about it. Hence, given a medical concept in a HP we transverse the medical ontology and collect the TD associated with it. In case a conceptual match is not possible, the ontology allows us to select a TD related to the specialisation or generalisation of the concept.

In Patel V, Rogers R and Haux R (Eds.) 10th World Congress on Medical Informatics (MedInfo'2001), September 2 – 5, London. IOS Press, Amsterdam.

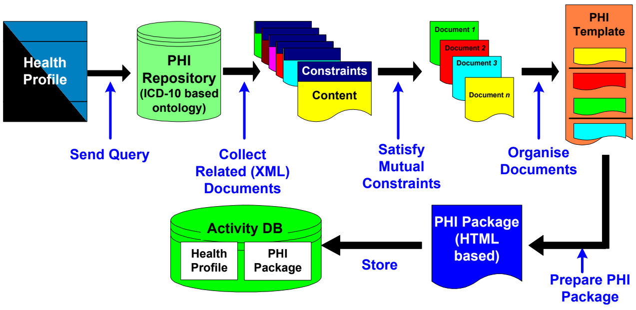

Figure 2 - The process flow for PHI composition

(ii) The collected TD are presented to a *Constraint Satisfaction Engine*—employing constraint logic programming techniques—to ensure the medical consistency of the multiple, heterogeneous TD when aggregated to yield a seamless PHI package.

Note that each TD contains a set of local constraints that need to be satisfied for it to be included in the final PHI package. For instance, if we have collected two TD, one TD about lifestyle recommends a particular diet which might have a high sugar content, whereas the other TD about diabetes maintenance prohibits the use of sugar-based substances. In this case, the two TD give recommendations that contradict each other. The constraint satisfaction engine is responsible for satisfying such constraints by simultaneously operating on the local constraints of all TD using medical knowledge (rules) stored in the PHIDS medical knowledge base. At the conclusion of the constraint satisfaction process, the mutually inconsistent TD are filtered out, whilst the mutually consistent TD are selected to compose the PHI package as per the specification of the PHI template.

(iii) The selected TD need to be seamlessly aggregated to yield a PHI package that appears as a continuous, structured and readable PHI document. To ensure a standard feel-and-look of the PHI package, we use a *PHI Template*—a specification that defines how the disparate pieces of information (i.e. multiple TD) are to be aggregated. Technically, the PHI template is a XML style-sheet (XSL)—a pre-defined structure defining place-holders for imported text (in XML format) and graphics (see Figure 3).

(iv) The PHI document is derived by placing XML-based TD at the designated places in the PHI template.

(v) Finally, the XML-based PHI package is converted to HTML format for Internet based delivery. Figure 4 shows an exemplar HTML-based PHI package, comprising the PHI content, the banners at the top and the dynamically created navigation frame at the left of the HTML page.

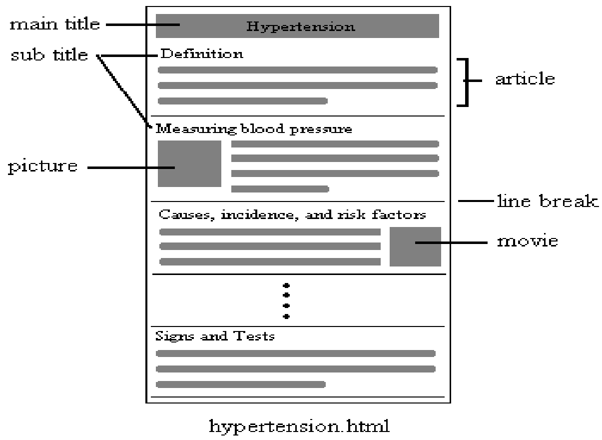

Figure 3 - An Exemplar PHI template

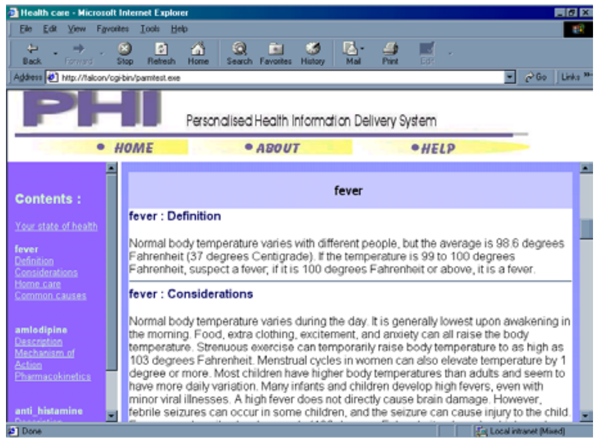

Figure 4 - Screenshot of the PHI package describing *Fever*.

Task 3: Delivery of the PHI package

PHIDS incorporates both *Push (System-Motivated)* and *Pull (Client-Motivated)* modes for information delivery. The operational characteristics are different for these two delivery modalities as explained below:

Client-Motivated Mode involves the typical *Pulling* of desired information from a web-site. In this case, the user requests for PHI, which is then composed and sent to the user's web browser (as shown in Figure 5).

System-Motivated Mode is the featured and innovative delivery mode whereby up-to-date PHI is pro-actively and periodically *Pushed* to users, over the Internet to their email accounts—i.e. just-in-time PHI automatically delivered at the desktop. For registered users, PHIDS takes charge of their dynamic health needs and pro-actively compiles and delivers the 'best' PHI at scheduled intervals. We argue that by way of the implementation of the push delivery mode, PHIDS can be perceived as a *pro-active health guardian*.

The operational functionality of the Push mode is made feasible by the implementation of a back-end *Delivery Manager*—akin to an autonomous intelligent agent—that monitors events that prompt the pro-active generation of PHI (see Figure 6). The Delivery Manager decides to push information, via email, under the following two conditions:

- Update/Addition of healthcare content that directly corresponds (or is even indirectly relevant) to the most recent HP of an individual. Note that the most recent health profile and the last version of the PHI package is stored in the *Activity Database*.
- A change in the health status of the individual is noted vis-à-vis an update to the individual's EMR We record the date of the EMR based on which the most recent HP is generated. If the date of the most recent EMR is later than the date of the most recent HP, then it implies that the current HP has 'aged'. The Delivery Manager then initiates the generation of a recent PHI package and subsequently 'pushes' it to the individual.

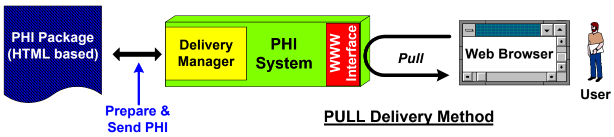

Figure 5 - The process flow for PULL based PHI delivery

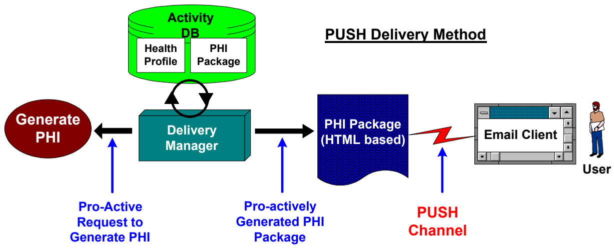

Figure 6 - The process flow for PUSH based PHI delivery

An important feature of the Delivery Manager is that it prevents the repetitive delivery of previously-sent information. This is achieved by comparing the contents of the new PHI package with the previous 2-3 PHI packages stored in the Activity DB. We believe that the repeated presentation of the same information might diminish the interest of the user.

Using Constraint Satisfaction Techniques To Compose PHI

We mentioned earlier that composition of the PHI involves the selection of multiple TD from a pre-defined PHI repository (comprising a large number of TD), such that the selected TD can co-exist in the same PHI package without violating any medical constraints. This is achieved by a constraint satisfaction engine implemented within PHIDS.

Typically, a *constraint* describes a relation between components (represented as variables) and all the allowed combinations of values that can be assigned to the components as per the stated relation. In the context of PHI composition, constraint satisfaction is an *abductive reasoning* task whereby out of the set of all possible combinations of TD (corresponding to the individual's HP), we select only those TD which satisfy pre-specified medical constraints—the constraints determine that information-item *x* and *y* do not go together, hence ruling out various 'inconsistent' combinations of TD.

In our framework, a constraint is defined as a tuple:

$$constraint\_type([(component_1, weight_1), (component_2, weight_2), \ldots, (component_n, weight_n)])$$

for instance

$$constraint\_drugs_A([(drugX, 0), (drugY, 0.5)]) \quad (1)$$
$$constraint\_drugs_B([(drugX, 1), (drugZ, 0.7)]) \quad (2)$$
$$constraint\_lifestyle_A([(rest, 0), (fatty\_diet, 1)]) \quad (3)$$
$$constraint\_lifestyle_B([(rest, 0)]) \quad (4)$$

The weight associated with each constraint component indicates its *degree of recommendation*—i.e. a weight of value 0 implies that the component is recommended, whereas a weight value of 1 implies that the component is not recommended. For example, constraint no. 3 (shown above) recommends that the patient rests and at the same time not recommends a fatty diet. The weight values are continuous, ranging from *recommended* to *not recommended* i.e. from 0 to 1, with a value of 0.5 indicating a *neutral* stance towards the constraint component.

We indicated earlier that each TD comprises a number of constraints and during PHI composition we compare the constraints of multiple TD to select only those TD that manifest globally consistent constraints. For instance, constraints of TD_A will be compared with the corresponding constraints (of matching type) of TD_B. Constraint satisfaction takes place according to two methods:

Constraint Satisfaction : Method I

In the simplest form of constraint satisfaction we compare similar constraint components from two different TD. The weights of the constraint components are subjected to the following formula:

$$0 \leq \left| WeightX_{sub-doc1} - WeightX_{sub-doc2} \right| \leq 1$$

whereby the output value 1 indicates a severe conflict, whilst an output of 0 implies no conflict between the information content of the two TD. Consider the following two constraints:

$$drug_{Antihistamine}\ [(aspirin, 0.2), (panadol, 0)] \quad (TD_A)$$
$$drug_{Metaformin}\ [(panadol, 1), acetaminop\ hen, 0] \quad (TD_B)$$

Note that TD_A strongly recommends Panadol to be taken together with Antihistamines, whilst TD_B strongly not recommends the use of Panadol with Metaformin. Henceforth, logically these two TD cannot co-exist together within the same PHI. The constraint satisfaction mechanism concludes the same as shown below:

$$\left| Weight_{Panadol}^{Metafor\min} - Weight_{Panadol}^{Antihista\min e} \right| == 1$$

implying that either of TD_A or TD_B will be excluded from the final PHI, depending on *the total conformance score*—a user defined parameter indicting the conformance of the TD with all selected TD—of each TD.

Constraint Satisfaction – Method II

This is a more advance method of constraint satisfaction, involving a constraint database (compiled by medical experts) that comprises binary constraints defining a compatibility relation between two constraint components. Below we show exemplar constraints between (a) drugs and treatment and (b) drugs and lifestyle:

$C_{Drugs,\ Treatment}$ = *(drugX, treatmentY, compatibilityXY)*

$C_{Drugs,\ Lifestyle}$ = *(drugA, lifestyleB, compatibilityAB)*

The variable *compatibility* determines the measure of compatibility between the two constraint components indicated in the constraint. The compatibility measure ranges from –1 (i.e. the constraint components are highly not compatible) to 1 (i.e. the constraint components are highly compatible).

Given two TD with the following constraints:

drugs[(drugX, weightX), (drugY, weighty)] (for TD_A)

lifestyle[(lifestyleA, weightA)] (for TD_B)

Constraint satisfaction between two TD is carried out by firstly combining the constraints to form a more complex constraint of the form:

$C_{Drugs,\ Lifestyle}$ = [(drugX, lifestyleA, (weightX-weightA)), (drugY, lifestyleA, (weightY-weightA))]

Each individual constraint in the above combined constraint is checked against the constraints stored in the constraint database. The compatibility measure between the constraint components *drugX* and *lifestyleY* is first determined from the constraint database. Next, the below given formula is applied to determine their compatibility with each other.

$$-1 \le \left(\left|WeightX - WeightK\right|\right) - CompatibilityXK \le 1$$

If $\left(\left|WeightX - WeightK\right|\right) - CompatibilityXK$ == 1

then it implies that the two constraint components cannot co-exist together, hence one of the TD need to be discarded.

The above two constraint satisfaction methods form the basis for the dynamic composition of PHI.

Concluding Remarks

Our work was initially motivated by the Malaysian Multimedia Super Corridor (MSC) Telemedicine initiative, which articulated the necessity for PHI dissemination to empower individuals to take charge of their day-to-day health promotion and preservation needs [2]. PHIDS can be regarded as a functional prototype for the aforementioned Tele-Medicine project, though with limited community and content coverage.

Currently, PHIDS is under trial, involving around 200 registered members whose EMR's are locally available within our institution. Soon we plan to initiate randomised trials to ascertain the efficacy of the system and will report the results in a separate publication.

References

[1] Abidi SSR and Yusoff Z. Telemedicine in the Malaysian Multimedia Super Corridor: Towards Personalised Lifetime Health Plans. In: Kokol P et al eds. *Medical Informatics in Europe'99*. Amsterdam: IOS Press, 1999.

[2] Gillespie M and Ellis L. Computer-based Patient Education Revisited'. *Journal of Medical Systems* 1993: 17(3-4).

[3] McRoy S, Liu-Perez A and Ali S. Interactive Computerized Health Care Education. *Journal of the American Medical Informatics Association* 1998: 5(4), pp. 347-356.

[4] Marine S, Guard R and Morris T. A Model for Enhancing Worldwide Personal Health and Wellness. In: Cesnik B et al eds. *Medical Informatics in Europe'98*. Amsterdam: IOS Press, 1998.

[5] Ball M. Medical Informatics in the New Millennium. In: Cesnik B et al eds. *Medical Informatics in Europe'98*. Amsterdam: IOS Press, 1998.

[6] Abidi SSR and Goh A. Intelligent Healthcare Information Dissemination Featuring Electronic Medical Record Profiling and Customised Information Composition. *Intl. ICSC Congress on Intelligent Systems and Applications*, Sydney, 2000.

[7] Hirst G et al. Authoring and Generating Health-education Documents that are Tailored to the Needs of the Individual Patient. In: Jameson A, Paris C, and Tasso C eds. *Proc. of the 6th Intl. Conf. on User Modeling*. 1998, Springer Verlag, pp. 107-118.

[8] Bental D, Cawsey A and Jones R. Patient Information Systems that Tailor to the Individual. *Patient Education and Counselling* 1999:36(2), pp. 171-180.

[9] Cawsey A, Bental D, Jones R, Pearson J and Carter E. A Personalised Patient Information System Using GRAIL. In: Richards B eds. *Current Perspectives in Healthcare Computing*. 1998, pp. 115-121.

[10] Strecher VJ, Kreuter M. The Effects of Computer Tailored Smoking Cessation Messages in Family Practice Settings. *Journal of Family Practice* 1994: 39(3), pp. 43-49.

Address for correspondence

Health Informatics Research Group, School of Computer Sciences, Universiti Sains Malaysia, 11800 Penang, Malaysia. Email: sraza@cs.usm.my