In 14th IEEE Symposium on Computer Based Medical Systems (CBMS'2001), 26-27 July, Bethesda (USA)

An Intelligent Info-Structure for Composing and Pushing Personalised Healthcare Information 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, 11800 Penang, Malaysia
sraza@cs.usm.my

Abstract

The featured work aims to provide a technology-enriched solution to web-mediated patient empowerment initiatives via the implementation of an intelligent info-structure that features (1) dynamic composition of Personalized Healthcare Information (PHI) conformant to an individual's EMR-based health profile; and (2) pro-active Internet-based delivery of PHI. Healthcare information personalization is achieved via profile-specific selection and template-specific aggregation of multiple fine-grained, pre-authored, generic 'topic-specific' healthcare information content. The use of intelligent constraint satisfaction techniques ensure the medical correctness of the dynamically composed PHI document. The application of Internet-based push technology allows for the PHI document to be pushed to the individual's email account, thus guaranteeing the timely availability of high-quality health maintenance information.

1. Introduction

Patient empowerment programs emphasize the dissemination of person-centric healthcare information content, designed with the objective to educate individuals to maintain a prolonged state of wellness. Lately, the proliferation of *Healthcare Portals* over the WWW have effectuated the ubiquitous dissemination of generic healthcare information applicable to a wide audience [1] [2] [3]. Notwithstanding the effectiveness of WWW-based patient empowerment initiatives, we believe that the impact of such initiatives can be further enhanced by providing the following value-added features [4]:

(I) *Personalized Healthcare Information* (PHI) conforming to an individual's current *Health Profile* (HP). This will ensure that the healthcare content disseminated to an individual is specifically focused towards his/her prevailing healthcare needs, akin to the kind of personalized service one enjoys from a visit to a medical practitioner.

(II) Pro-active delivery of PHI to the individual—i.e. the PHI should be periodically *pushed* to the individual to ensure the timely consumption of the current PHI.

The featured work aims to add value to web-mediated patient empowerment initiatives via the implementation of an intelligent info-structure for dynamically composing and pushing PHI over the Internet. We present the design and functional characteristics of a *Personalised Healthcare Information Delivery System* (PHIDS) that exhibits the following features: (a) Automated generation of an individual's current HP from information contained in his/her Electronic Medical Record (EMR); (b) Scheduled generation of PHI based on an individual's current HP; (c) Composition of a comprehensive PHI document by systematically synthesizing multiple healthcare *Topic-*

specific Documents (TD), where each TD corresponds to some healthcare aspect pertinent to an individual's HP; (d) Use of intelligent constraint satisfaction techniques to determine inter-TD compatibility within a PHIP—i.e. ensuring that the contents of constituent TD do not medically contradict each other; (e) Medical ontology based organization of health information; (g) Use of Extensible Markup Language (XML) to facilitate both the representation and Internet-mediated distribution of healthcare information; and (h) Pro-active 'push-based' delivery of PHI over the Internet [4] [5].

2. Our approach for healthcare information personalization

Personalization—in terms of both content and form—of generic healthcare information, as per an individual's HP, has been previously addressed as follows: (a) the creation of a master document comprising a large set of blocks of healthcare text (with annotations) which are selectively included or excluded for both content and form [6]; and (b) the use of a knowledge base to select and synthesize a variety of individual healthcare texts—where each text pertains to a specific healthcare issue—to yield a composite healthcare document. Here a knowledge base of rules is required to facilitate the selection of content, prepositional ordering and lexical choice [7].

Our approach for information personalization draws inspiration from the above two approaches. As per our approach, information personalization originates from the profile-specific selection and template-specific combination of multiple fine-grained, pre-authored, generic 'topic-specific' healthcare information content. In our framework, composition of PHI is carried out as a three-step process: (1) Selection of a set of 'highly-relevant' TD from a large pool of TD, where each selected TD specifically addresses some healthcare concern reported in the individual's HP; (2) Combination of the selected TD to yield a composite PHI document; and (3) Verification of the medical correctness of the cumulative PHI document—i.e. ensuring that the constituent TD do not suggest information/recommendations that are medically inconsistent.

In a real-life operational situation, there exist the possibility that two (or more) TD may contain information that is medically inconsistent if presented together. The novelty of our information personalization approach derives from the use of intelligent constraint satisfaction techniques to determine inter-TD compatibility—i.e. we satisfy the medical constraints associated with each selected TD in order to determine what information to exclude from what to include within the eventual PHI document.

3. 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.
2. *Composition of a PHI Document* by systematically amalgamating multiple TD based on the individual's most current HP.
3. *Delivery of the PHI Document* using both pull (client-mediated) and push (system-mediated) mechanisms.

3.1. Health profile generation

The generation of an up-to-date individual-specific HP demands the collection and subsequent summarization of relevant and most recent information from the individual's EMR. Functionally, HP generation is carried out as follows (see Figure 1):

(i) The temporal span of the HP is determined by a *Health Profile Window*—the retrospective timeframe spanning from the current date to some specified past date (typically set to 3 weeks) during which the EMR is examined for information collection.

(ii) The HP is generated by selecting certain relevant fields of the EMR—i.e. fields that refer to (a) acute diseases and (b) episodic encounters with healthcare providers.

(iii) The HP content is standardised with respect native terminological standards. The *Information Standardizer* 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) If deemed necessary, 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 document. For future reference purposes, the final HP is stored in the *Activity DB*.

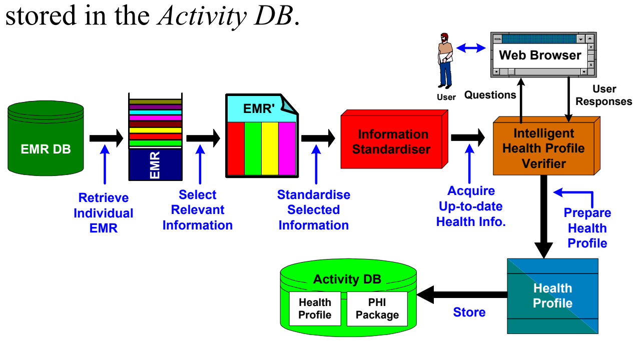

Figure 1. The process flow for HP generation

3.2. Composition of the PHI document

The PHI being composed comprises three main sections: (1) Information about long term-term clinical conditions and management regimes; (2) Information about short-term non-acute illness episodes and associated therapy and rehabilitation; and (3) General healthcare education with a wellness maintenance connotation. Each PHI section manifests a systematic aggregation of corresponding TD, such that each individual illness/concern/issue noted in the HP is addressed at least by a single TD. The information content and structure of an exemplar TD is shown in Figure 2. Note that the upper portion of the TD contains standard information applicable to all user profiles. But, the bottom section, containing medically-valid and focused suggestions, is personalized by the selective inclusion of profile-specific information derived from a set of specialized *Sub Topic- Specific Documents* (STD) pre-associated with the TD.

We mentioned earlier that intelligent constraint satisfaction techniques are used to ensure that the aggregated PHI document is medically consistent—i.e. upon aggregation the individual TD do not contradict each other or lead to improper recommendations. Each TD contains a set of *standard constraints* applicable to its standard information section. Similarly each STD, associated with a TD, also contain a set of *personalization constraints* that are local it. The global constraints of a TD are a combination of its standard constraints and the personalization constraints of all constituent STD's.

In our framework, a *constraint* describes a relation between components 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. Typically, a constraint is defined as a tuple (see Table 1):

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

Table 1. Exemplar constraints for three independent TD's

TD_A	TD_B	TD_C
$drugs_A$([(drugX, 0), (drugY, 1)]) $lifestyle_A$([(rest,0), (fatty_diet, 1)])	$drugs_A$([(drugX, 1), (drugY, 0.7)]) $allergy_A$([(allergyX, 0)])	$drugs_A$([(drugX, 0), (drugZ, 0), drugA, 1, drugB, 0.5])

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, the lifestyle constraint for TD_A (shown in Table 1) recommends that the patient rests and at the same time not recommends a fatty diet. Consider a simple illustration of constraint satisfaction: Given the set of constraints for TD_A, TD_B and TD_C in Table 1, the constraint satisfaction algorithm will conclude that TD_A and TD_B are incompatible with each other because, both TD_A and TD_B contain a constraint that *recommends* and *not-recommends* drugX, respectively. Logically speaking we cannot recommend and not-recommend something at the same time, hence one of the TD's will be discarded, depending on the *cumulative compatibility score* of the TD. Further note that TD_A and TD_C are compatible with each other as they both recommend drugX and their other constraints also do not contradict each other.

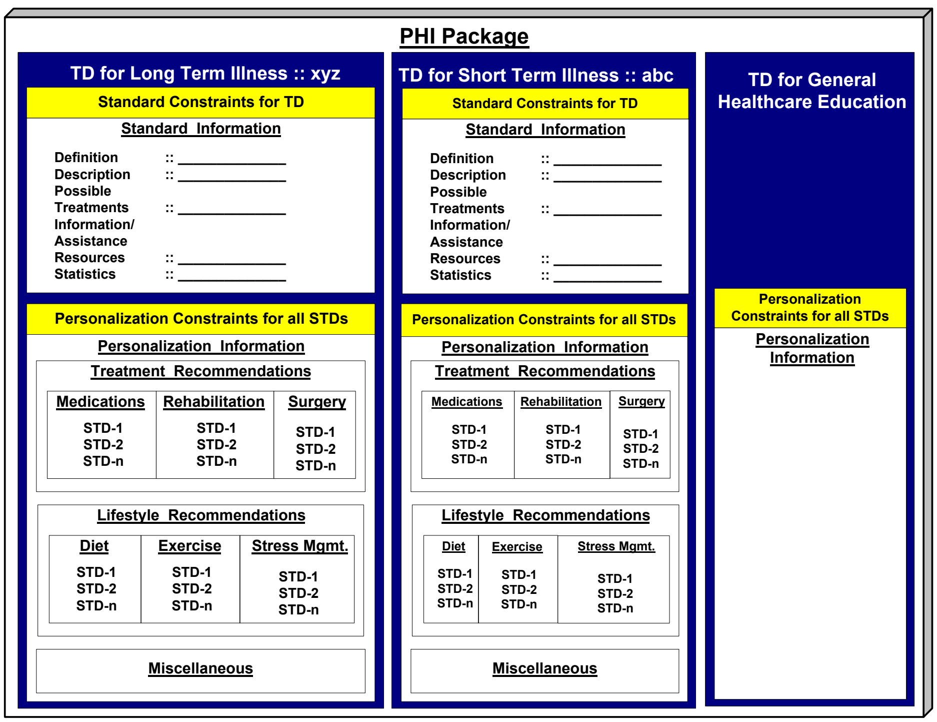

Figure 2. The information content and structure of a TD

We briefly discuss the process for PHI composition (as shown in Figure 3) 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 Healthcare Information Repository. Each TD is represented as an XML document.

(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 document.

(iii) The selected TD are aggregated according to a *PHI template*—a specification defining place-holders for imported text (in XML format) and graphics—to yield a continuous, structured and readable PHI document.

(iv) Finally, the XML-based PHI document is converted to HTML format for Internet based delivery.

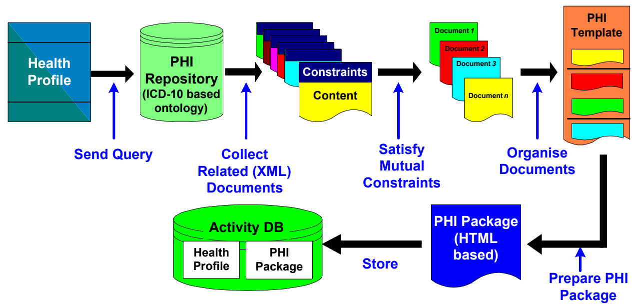

Figure 3. The process flow for PHI composition

3.3. Delivery of the PHI document

PHIDS incorporates both *Push (System-Motivated)* and *Pull (Client-Motivated)* modes for information delivery. Their operational characteristics are 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 dynamically composed and sent to the user's web browser.

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. For registered users, PHIDS takes charge of their dynamic health needs and pro-actively compiles and delivers the 'best' PHI at scheduled intervals. The operational functionality of the Push mode is made feasible by a back-end *Delivery Manager* that prompts the pro-active generation of PHI. 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.

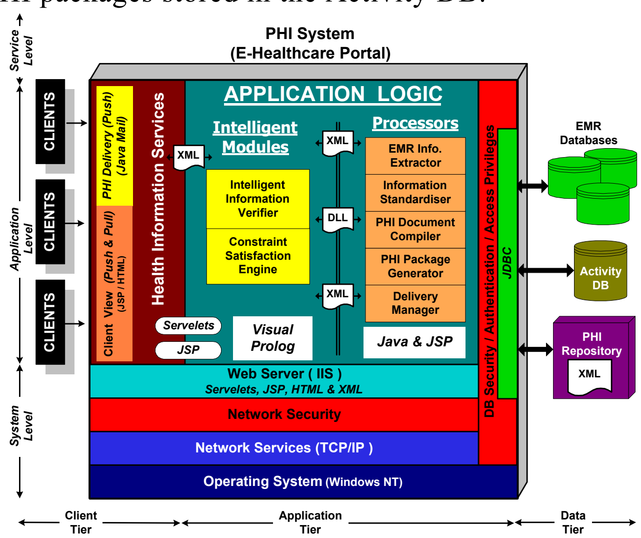

Figure 4. System architecture of PHIDS.

4. PHIDS implementation platform details

PHIDS is as an Intelligent Info-Structure with its intelligent core implemented using the Prolog programming language, specifically Visual Prolog version 5. Java and Javascript is used for the implementation of a suite of processors and also for all web-based interfaces for information dissemination and collection. Healthcare information content is developed in XML format, with standardized XSL for style and DTD for content specification. MS SQL server is used for all database-related activities. Finally, we use Microsoft's IIS server as the web server for the delivery of PHI. Figure 4 shows the multi-tiered architecture of PHIDS, illustrating the various modules.

5. Concluding remarks

The effectiveness of patient empowerment initiatives can be quantified by the 'improved' ability of individuals to make 'informed' judgments about personal health maintenance. This necessitates a re-design of traditional information dissemination mechanisms; for maximum impact it is imperative to effectuate the delivery of *right* healthcare information for consumption at the *right* instance. The featured work presents a demonstrator application—with the potential for large-scale deployment—that provides a novel outlook to (a) healthcare information personalization via the use of constraint satisfaction techniques, leveraging a substantial volume of medically-valid constraints and healthcare information content authored by experienced medical practitioners; and (b) web-mediated information delivery, employing both *push* (system-initiated) and *pull* (user-initiated) oriented mechanisms.

At present we are working to extend the functionality of the said system via the (1) application of medical ontologies for content classification and compilation of PHI document; (2) ensuring content privacy during various healthcare-related transactions, and (3) the use of intelligent technologies, in particular constraint logic programming, to satisfy complex inter-dependencies arising during dynamic text composition. PHIDS is under trail and a study of its effectiveness will be released as a publication later.

References

[1] M. Gillespie and L. Ellis, "Computer-based Patient Education Revisited", Journal of Medical Systems, 1993.

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

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

[4] S.S.R. Abidi and A. Goh, "A Personalised Healthcare Information Delivery System: Pushing Customised Healthcare Information Over the WWW". In: Hasman A, Blobel B, et al eds. Proc. of Medical Informatics in Europe'2000, Amsterdam, IOS Press, 2000.

[5] S.S.R. Abidi and A. Goh, "Intelligent Healthcare Information Dissemination Featuring Electronic Medical Record Profiling and Customised Information Composition", International ICSC Congress on Intelligent Systems and Applications, Sydney, 2000.

[6] G. Hirst, C. DiMarco and L. Wanner, "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, Springer Verlag, 1998, pp. 107-118.

[7] V.J. Strecher and M. Kreuter, "The Effects of Computer Tailored Smoking Cessation Messages in Family Practice Settings", Journal of Family Practice, 39(3), 1994, pp. 43-49.