Using Computerized Clinical Practice Guidelines to Generate Tailored Patient Education Materials

Brent Jones, Syed Sibte Raza Abidi, Winston Ying
Health Informatics Lab, Faculty of Computer Science, Dalhousie University, Halifax, Canada
jonesbd@dal.ca & sraza@cs.dal.ca

Abstract

In this paper we present a novel approach to generate tailored patient education materials by using an evidence-based clinical practice guideline to model patients and direct the selection of relevant education content. We converted a Canadian guideline for managing dyslipidemia into the XML-based GEM formalism, and then modified it to support the tailoring of patient education material. We created the information content suitable for tailored messages, along with a document structure needed to present them. A rules engine was developed to process the newly converted GEM-based guideline and create the tailored documents based on patient data entered.

1. Introduction

Clinical practice guidelines (CPGs) are documents written for physicians that provide recommendations about the investigation and management of specific health problems. In general, CPGs over the last 10 years have been developed based on reviews of medical literature to determine the best evidence for the recommendations made. Developed from the most current scientific evidence to date, CPGs represent a solid knowledge resource for physicians, to help them provide care to their patients. Patient education is a crucial part of managing many health problems, particularly chronic diseases. A great number of long-term illnesses are influenced by health behaviours, such as dietary habits and exercise. Changing or modifying harmful behaviours can help to prevent the onset or progression of chronic disease and patient education is felt to be an important strategy to help patients achieve this. Print materials, such as handouts, and face-to-face counselling by physicians are examples of patient-directed education initiatives.

Despite the wide availability of a great number of CPGs there is a lack of connection between the recommendations they make and the actual clinical care provided to patients. Physicians have identified a number of barriers to implementing guidelines and providing the patient education important for behaviour change. Of the different strategies that have been tested to overcome these barriers, a combination of reminders and patient-directed interventions, such as patient education, shows promise. Tailoring education materials to patients has a strong potential to increase their efficacy.

We believe that a CPG can serve not only as a decision support reminder tool for physicians, but its inherent decision logic can be leveraged to create patient profiles that can subsequently form the basis for tailoring patient-specific educational intervention materials. CPGs by design entail a decision logic that is structured in an algorithmic format intended to support clinical decision making by practitioners. We argue that the same decision logic in a CPG can be used to profile patients for generating tailored information. Simply put, we are suggesting that the recommendations or actions at specific decision or choice points in a CPG can be supplemented with patient education content related to both the suggested action and the conditions that led to that particular action. As a result, this approach to patient profiling is driven by the CPG's decision logic, which is based on the best current clinical evidence.

In this paper we will present our approach for converting a paper-based CPG to GEM (the Guideline Elements Model) format, with the intention of using the GEM-based CPG as the basis for generating tailored education materials. This is a unique and interesting utility of CPGs as they are not designed with this purpose in mind. CPGs were developed to provide physicians a framework to manage specific health problems in their patients. Our approach takes this framework and uses it to create relevant information materials for patients. It involves a series of important steps, including converting the CPG into a computer-interpretable format, developing the relevant education materials, identifying the logic for presenting those materials and creating an approach to implementing the system. These steps and their associated challenges will be addressed in this paper.

2. Background

This research was pursued to bridge two important issues in clinical medicine: patient education and implementation of CPGs as recommender tools. Our application works on both fronts, providing tailored education materials to patients and computer decision support (CDS) to physicians. The key to our system is that is uses an executable version of a CPG as the knowledge base for both components. In this way, all output it generates is based on clinical evidence.

2.1 Clinical practice guidelines

Traditionally, clinical experience, epitomized by the "expert" or panel of experts, along with an understanding of the basic mechanisms of disease formed the knowledge base of medicine [1]. Conversely, evidence-based medicine (EBM) is the conscientious and judicious use of current best evidence from clinical care research to guide health care decisions [2]. An underlying belief of EBM is that experts are more fallible in their recommendations of what is successful for treating patients than evidence based on rigorous health care research [1]. EBM developed as a way to base decisions about patient care on quality evidence from applied research in complex clinical settings rather than solely on experience and an understanding of basic disease principles.

Clinical practice guidelines developed in the last 10 years can be thought of as a tool for implementing the principles of evidence-based medicine. They are systematically developed statements to assist both practitioner and patient decisions about appropriate health care for specific clinical circumstances [3]. CPGs typically address a specific health condition, such as dyslipidemia, and provide recommendations to the physician about issues such as who to investigate for the problem, how to investigate it, how to diagnose it, and how to treat it. While CPGs are not a new concept, they are only recently being developed based on principles of EBM, the systematic evaluation of clinical research [4]. The main goals of CPGs are to improve outcomes for patients, assist clinical decision making, reduce unnecessary costs and diminish inappropriate practice [3,5,6].

EBM and CPGs are not without their criticisms, though. Issues such as increased costs to fully implement EBM and the difficulty of applying results of clinical trials to individuals have been raised [1]. Perhaps the biggest objection is the perception that CPGs promote a "cookbook" approach to providing care, or treating all patients the same according to a formula or algorithm derived from research [1].

In response to the latter criticism, Haynes [1] points out that this formulaic approach was never the intention. EBM should be considered an adjunct to traditional care, used in conjunction with other components of decision-making, including the patient's circumstances (as assessed through the expertise of the clinician) and the patient's individual preferences. Clinical judgement and expertise should always be considered essential to the decision process.

2.2 The treatment gap

Despite the great number of CPGs and their wide dissemination in journals and over the Internet, there is a difference between the best practice (based on scientific evidence) and actual clinical care. Grol et al [7] report that at least 30-40% of patients do not receive care according to the current scientific evidence and at least 20% of the care provided is unwarranted or potentially harmful. This difference between optimal practice and actual care has been labelled the "treatment gap" [8]. A number of theories try to explain the treatment gap. Some point to the guidelines themselves, for example, referring to specific attributes such as the type of problem addressed or the degree of complexity of the recommendations. Other theories focus on barriers related to the health care system and its stakeholders. Examples include financial disincentives, organizational constraints, standard practices, opinions of colleagues, medical training, information overload, clinical uncertainty and self confidence [9,10,11].

The literature is clear that a single CPG implementation intervention is less effective than a multifaceted approach to change involving two or more strategies [9,12,13]. Two strategies that have been associated with improvements in performance are reminders, such as computerized decision support, and patient-directed interventions, such as patient education [14].

2.3 From CPG to computerized decision support

A computerized decision support system (CDS) is an example of an implementation intervention based on reminders. Reminders are a type of feedback whose goals are to replace memory and to inform decisions with useful, timely, relevant information [9]. They have been shown to be the most effective of all the implementation strategies geared toward the practitioner [11]. CDS is a reminder system that compares patient characteristics with a knowledge base and then guides a health provider by offering patient specific and situation specific advice [15]. The knowledge base forms the rules of a particular clinical situation to which the patient's relevant data can be applied. A review of the CDS literature by Hunt et al [16] reveals that 48 of 68 controlled trials of CDS systems showed benefits. Systems targeting drug dosing, preventative care and "other medical care" (a variety of conditions and outcomes, primarily assessing the process of care) showed the greatest effects.

A number of modeling methodologies have been developed to convert natural language guidelines into an electronic format that is executable by a computer, thus allowing the CPG to become a CDS system. Examples of such methodologies include GEM, EON, GLIF, GUIDE and Prodigy [17]. For our research we worked with GEM, the Guideline Elements Model [18].

GEM was developed at the Yale Center for Medical Informatics and was designed to provide a structure for marking up any CPG in XML. It is essentially a DTD comprised of more than 100 elements that describe all aspects of the document and places them into categories that define the CPG, according to its identity, developer, purpose, intended audience, target population, method of development, testing, review plan and knowledge components [18]. The knowledge components section is reserved for the clinical content, the focus of any CPG. Here, the content is broken down into one or more recommendations. Each recommendation is further tagged as either conditional or imperative. Conditional recommendations apply to situations specified by an "if-then" statement that acts on decision variables and results in an action. Each of these recommendations has one or more logic elements that are meant to hold the "if-then" statements. Imperative recommendations are directed at the entire CPG target population without restriction. Each of these recommendations results in a directive.

To facilitate more consistent conversion of a CPG to XML, the developers of GEM created GEM Cutter [19,20], an authoring tool that allows the person doing the conversion to simply copy the relevant text of a CPG (in .txt or .rtf format) and then paste it into the pertinent GEM element. Once the natural language of the CPG has been converted to a GEM document (marked up in XML) it becomes possible for a computer to process the instructions within.

In this way, the wealth of evidence-based knowledge in the original CPG can be implemented using a strategy that has been shown to have positive effects on the adoption of CPGs by physicians and on the process of care. Furthermore, by having the CPG in a format that can be processed by a computer, it is also possible to have the CPG do something not originally intended: it can act as the basis for computer-generated tailored patient education materials. The process for doing this with a GEM-based guideline needs to be defined and this is the focus of our research.

2.4 Patient education

Patient-directed interventions are strategies designed to improve clinical practice that focus directly on patients [11]. One example is patient education, or any set of planned, educational activities designed to improve patients' health behaviours and/or health status [21]. In most cases, the goals of patient education are to maintain and improve health, but in some cases they may be to slow deterioration [21]. These goals of patient education are addressed by influencing changes in patients' behaviours and/or mental attitudes [21].

Patients today are more responsible for managing their own illnesses [21]. Whereas physicians are primarily responsible for the medical management of the disease, the patient is primarily responsible for the day-to-day management of the illness, the consequence of the disease [21]. This is particularly true for chronic conditions such as diabetes or hypertension where much of the management relies on patients' health behaviours, especially diet and exercise. The physician can make recommendations or write prescriptions but it is the patient who must then follow through and implement the instructions and modify behaviours. For patients to be able to do this they need to acquire certain information and skills (and the confidence to apply those skills), as well as the ability to cope with changing roles and emotions [21]. Fortunately, patients want to participate in decisions about their health care and their desire for information is high [22]. What is more, there is support in the literature that patient-directed interventions, such as patient education, are important and beneficial to care [14,23,24]. A review by Grimshaw et al [14] showed that patient-directed interventions might result in moderate to large improvements in performance, particularly related to prevention of disease and Tang and Newcomb [22] found that providing education materials to patients can improve their understanding and satisfaction, and can increase their motivation to adhere to treatment plans. Self-management education makes the patient more of an active partner with the physician in managing health problems and less of a passive consumer. The key to making this partnership work is communication [21], and one accessible way to communicate the message is through print materials.

Two recent studies examined what patients want with respect to information from their physicians [22]. They focus primarily on print materials, such as pamphlets and handouts. The initial study revealed that patients want more information about their illness and treatment plan than they typically receive during physician visits. They also want information that is custom-tailored to their own situation; information when they formulate questions, which is generally after leaving the clinic and not during the physician visit; information that is endorsed by their physician as credible and applicable to their specific problem; information from other sources, such as journal articles and websites; and, information that they can retain for future reference. The second study reinforced these findings, particularly the desire for a personalized record.

According to Petty and Cacioppo's Elaboration Likelihood Model (ELM), people are more likely to actively and thoughtfully process information when they are motivated and able to do so, and perceived personal relevance of the message is one of the most important determinants of motivation [25]. Studies have shown that messages that are actively processed (considered carefully, related to other information encountered, compared to past experiences) are more likely to lead to enduring attitude and behaviour change [25].

Since patients want personalized information and it has been shown to be more likely to result in behaviour change there is a strong rationale for this approach. Lately, it has become possible to use computers to provide such personalized education materials to patients through what is aptly termed computer-tailoring. With computer-tailoring, the diagnostic and intervention expertise of physicians and other health educators is documented in a computerized system that allows the production and distribution of personalized education materials [26].

2.5.1 Tailored patient education. With respect to patient education, tailoring refers to any combination of information or change strategies intended to reach a specific person, based on characteristics that are unique to that person, related to the outcome of interest, and have been derived from an individual assessment [27]. In other words, health information can be tailored to fit a single person based on a profile created from a range of variables specific to that person.

The process of tailoring information materials can occur on a number of levels. Adding patients' names to documents is perhaps the simplest method to adapt materials to an individual and this is referred to as personalizing information [26,27]. Tailoring is much more than this, though. Information materials can be tailored to an individual by adapting the content of the health message, the desired change strategy, the source of the information provided, the method by which it is disseminated and the manner in which it is displayed [24].

How information is tailored to an individual can be based on a number of different variables. Sociodemographic characteristics are perhaps the most apparent and include determinants such as age, gender, race and socioeconomic status. Physical variables about an individual can also be used to adapt materials. Examples include a person's medical diagnosis, clinical examination findings and lab test results. A number of other variables derived from various behaviour theories can also be used. For example, Brug et al [27] describe behavioural, motivational and psychosocial factors from social cognitive models of behaviour. These factors include an individual's attitudes, social influences, personal norms, perceived behaviour control and intentions.

According to de Vries et al [26], computer-tailoring requires at least four components. These include: 1) a profile of the patient, or a diagnosis at the individual level of characteristics relevant to a person's health behaviour or illness; 2) a message library containing all the health education messages that might be needed; 3) an algorithm, or a set of decision rules that evaluates the diagnosis and generates the appropriate messages; and, 4) a channel or medium to deliver the message to the intended user.

There is support in the literature that tailored information materials outperform standard health education messages [27], and compared to non-tailored materials, tailored ones are more likely to be read and remembered, rated as attention catching, saved and discussed with others, and perceived as personally relevant [25].

3. Using CPG to tailor patient education material: Our approach

The purpose of this research was to demonstrate how a CPG converted to a GEM document can be used as the basis for generating tailored patient education materials. This approach creates a computer system that is built on two effective strategies for implementing CPGs—i.e. reminders and patient-directed interventions. Furthermore, our approach creates education materials that are tailored to patients based on profiles derived from CPGs. By using a GEM-based CPG as the cornerstone of the application, our approach creates profiles of patients based on the best current clinical evidence. These profiles are used to direct the selection of education materials specific to the patient.

The tasks required to achieve this can be divided into two major categories. The first relates to the conversion of the CPG into XML so that it can be processed by a decision support system for physicians. The second category of tasks relates to the selection and presentation of information to yield patient-specific education materials. These two tasks render the computerized CPG into a computer tailoring system for patient education.

3.1 GEM conversion

The CPG chosen for this project was Recommendations for the Management of Dyslipidemia and the Prevention of Cardiovascular Disease: 2003 update [28]. This is a natural language document that was distributed in paper and electronic (.pdf) formats in the Canadian Medical Association Journal. This guideline was chosen because of its evidence-based approach, its comprehensiveness and because of the significance of cardiovascular disease in North America. It focuses on the investigation and treatment of dyslipidemia as well as other risk factors for cardiovascular disease.

In preparation of converting this CPG to GEM, the original document was reviewed to identify each of the recommendations contained within. This process highlighted the fact that a document written in natural language can have a certain degree of ambiguity. For example, certain terms, such as moderate hypertriglyceridemia, were not explicitly defined in the

CPG. For the purposes of this study, personal clinical experience and other medical literature were used to resolve these ambiguities. Careful consideration was also given to the manner in which abstract concepts, such as central obesity, would be modeled as this had significant implications on the structure of our GEM document, its usability and the type of data we could gather once it was implemented.

3.1.1 Decision variables and action statements. All recommendations in the guideline were classified according to the GEM DTD's recommendation sub-elements. Each was designated as either conditional or imperative and a total of 15 conditionals were identified. For each of these, the associated decision variable(s) and action statement(s) were noted. One imperative recommendation and its directive were identified.

GEM Cutter was then used to mark up the CPG in XML tags by inserting the relevant text in the guideline into the appropriate GEM elements according to the design described above. This made the initial conversion quite straight forward, but as we progressed through the project several issues came to light that required attention. When we were adapting the GEM document to an executable format, we realized that there were some limitations to GEM and GEM Cutter and some problems with the original conversion.

First, a GEM-based CPG is divided into recommendations. For each recommendation, one declares its decision variables and action statements. If a variable such as age is present in every recommendation in the guideline, according to the GEM DTD it is declared each time. This creates multiple occurrences of the same variable. GEM Cutter automatically adds a sequential unique identifying attribute number to each decision variable and action as it is created. Therefore, if the variable age is present more than once in the GEM document, each instance is assigned a different attribute number. It was important for our application that all instances of the same variable have the same attribute number so that values attached to them could be entered only once, not each time the XSL stylesheet encountered it. GEM Cutter cannot do this so the attribute numbers in the GEM document were changed using a generic XML editor. An alternative would be to simply create each decision variable once when doing the GEM conversion with GEM Cutter, and refer to that variable by its unique attribute number in subsequent logic statements in the guideline.

Second, the decision variables stipulated in the CPG are of different data types. Some have fixed categorical values, such as male or female, while others require continuous values, such as the LDL-C level. Our application dynamically creates input forms to capture patient data based on the decision variables defined in the GEM document. In order to create user-friendly forms, it was necessary to have some way to identify for the XSL stylesheet the type of data each decision variable represented. This way, decision variables with categorical values such as male and female could be presented with radio buttons, and text boxes could be used for recording continuous values. To address this issue, we used the "value" sub-element of the GEM DTD. For categorical variables, we listed each of the possible choices in a separate value sub-element and the input form was set up to display these options as radio buttons. For continuous variables, we chose to identify them with the label "absolute value" in the value sub-element and the form was set up to have text boxes to capture this data. Our guideline also contains one decision variable that is a continuous value derived from a calculation. Since the application does the calculation, it was necessary to identify this variable so the XSL stylesheet did not present a text box asking for the value. We did this by using the label "calculated value" in the value sub-element.

Third, for our dynamically created input forms, we needed a way to indicate to physicians what data was required from them. We achieved this by inserting informative queries into the "description" sub-element for each decision variable. For example, to obtain a patient's abdominal circumference we created the element:

<description> What is this patient's abdominal circumference in cm? <description>

Fourth, our guideline contains an algorithm to calculate an individual's 10-year risk of coronary artery disease (based on the Framingham Heart Study). The risk is displayed as a percentage that is derived from the sum of five risk points related to age, gender, total cholesterol level, HDL-C level, smoking habits and systolic blood pressure. For example, an individual receives 2 risk points for being male and having an HDL-C level greater than 1.04 mmol/L. These are added to risk points in the four other categories to arrive at the 10-year risk percent. This calculation is an imperative function of the guideline as the result helps to drive later recommendations. In the original CPG, this algorithm is displayed as a form which the physician is required to complete, but GEM is not designed to display this type of text. We addressed this issue by breaking the algorithm into a series of conditional recommendations. The factors described above became decision variables and the risk points for each statement became actions. The rules of the algorithm were converted into approximately 150 individual "if-then" logic statements. Another option would have been to hard code the algorithm but our goal was to process it from within the GEM document.

Finally, because there are many ways to mark up any given document in GEM [29,30], we needed to do several conversions of our CPG until it was in a format that suited our purposes as a CDS system. For example, we initially copied only concise fragments of sentences, such as "routinely screen", from the document into the action elements. Since these elements contain information in the guideline that will be displayed as reminders to physicians, we decided to include larger portions of natural language from the original document. This created much longer recommendations but now they had context, making them more coherent. Because the application uses strictly what is available to it, all statements need to be self-explanatory.

3.1.2 GEM decision logic. Once the decision variables and actions were created, the decision rules for each recommendation in the guideline were written in the appropriate logic elements. These "if-then" statements are intended to use Boolean operators to define how to combine the decision variables and actions in a given conditional recommendation. In the statements, decision variables and actions are identified by their unique attribute number. For example, a recommendation stating "men over the age of 40 should have their fasting cholesterol levels routinely screened" would appear as:

"IF ((dv1=male) AND (dv2>40)) THEN a1".

Since the default for our system is to display an action on screen as a recommendation, we had to develop a way to handle other functions, such as calculating a patient's 10-year risk percent. We did this by modifying the then side of the logic statements as in the following example:

"IF ((dv1=male) AND (dv22<1.04)) THEN SET dv13 TO dv13 +2".

The variable dv13 was created to hold the risk points. As the rules engine compares the patient's data with each of the logic statements in the calculation it eventually sets dv13 to the correct value and presents it as the 10-year risk percent. This modification was also used to declare any categorical values derived from patient data. For example, a person's risk category (low, moderate or high) is dependent on his or her 10-year risk percent. The then side of the statement was used to set the value of the risk category variable (dv35) as follows:

"IF (dv13 <10) THEN SET dv35 TO Low".

The solutions we have presented to these issues allowed us to process all of the recommendations from within the GEM document. An alternative strategy would have been to have a programmer hard code some or all of the recommendations in Java, for example, but this was not our goal. GEM was created to facilitate translation of guidelines into a computer-interpretable format without requiring programming knowledge [29]. Our GEM implementation system (see below) builds on this by being a "generic" application that can process any GEM document marked up as we have suggested. That is, there is no code in it that is specific to the dyslipidemia guideline used in this research program. A CPG dealing with asthma, for example, could be converted to GEM by a clinician and then processed with our application.

3.2 Information tailoring system

The following tasks and challenges relate to the development of the information tailoring system. We followed the information tailoring framework by de Vries et al [26], but we added a fifth component—a document structure for presenting the education materials.

3.2.1 Patient profile. The first component we required was the diagnosis, or more specifically, the profile upon which to tailor the information. The CPG clearly establishes the variables and their accepted values that are to be used to guide decisions on diagnosing and managing dyslipidemia. For example, an LDL-C level of greater than 2.5 mmol/L in a high risk patient is considered elevated. Since the current scientific evidence shows that those variables and values are important for directing decisions about managing the clinical problem, then they should also be valid for directing the selection of patient education materials that discuss that problem. The patient profile for our information tailoring system, then, can be created from the variables identified in the CPG and is, therefore, evidence-based. That profile can then be used to objectively direct the selection of the appropriate education materials to provide to a given patient. For our application, tailoring of the documents was done based on a total of 12 demographic and physical decision variables from the CPG. The demographic variables include gender and age. The physical variables are based on diagnoses, clinical findings, laboratory values, 10-year risk of coronary artery disease and risk category. Examples include various cholesterol levels, a patient's smoking status and family history of cardiovascular disease.

3.2.2 Message library. The second component we required was a message library containing the information materials the computer would access. For our library content, we created tailored information materials from documents already available to patients. Permission was obtained from the College of Family Physicians of Canada to use their patient handouts and from the Canadian Heart and Stroke Foundation to use their online patient information materials. These documents cover major topics such as cardiovascular disease, dyslipidemia, hypertension, diabetes, weight loss and quitting smoking.

Our goal was to tailor the materials as closely as possible to the patient profiles being generated by the implemented GEM-CPG. Simply giving an entire cholesterol document to a patient with an elevated triglyceride, for example, was not tailored to a fine enough degree. To address this, we broke each of the documents into concise, focused sections called "information packets". Each packet discusses a single clinical topic and some are specific to gender or age category as well. In some, association to gender or age group was further enhanced by adding graphics and images. In this way it was possible to tailor the information for a given patient to a finer level of granularity than simply providing the original, generic, pre-written handouts. The pre-written content, though, essentially determined the theoretical basis upon which we were tailoring. The documents we chose take the Health Belief Model (HBM) approach by providing information related to risk, consequences of the illness and benefits of treatment. HBM is one of the most widely used frameworks of health behaviour and has been used extensively in health education interventions [21].

3.2.3 Document structure. An overall document structure was developed to organize and present the selected individual information packets into a coherent handout for a given patient (fig. 1). This document structure has four basic sections: an introduction, a risk factors section, a lifestyle information section and a references section.

The introduction provides a brief summary of the document and explains how it was created. It contains information that personalizes the document, such as the patient's name and birth date, and provides his or her risk category for heart disease, as determined by the CDS system. These data specific to each person are added dynamically to the introduction paragraph.

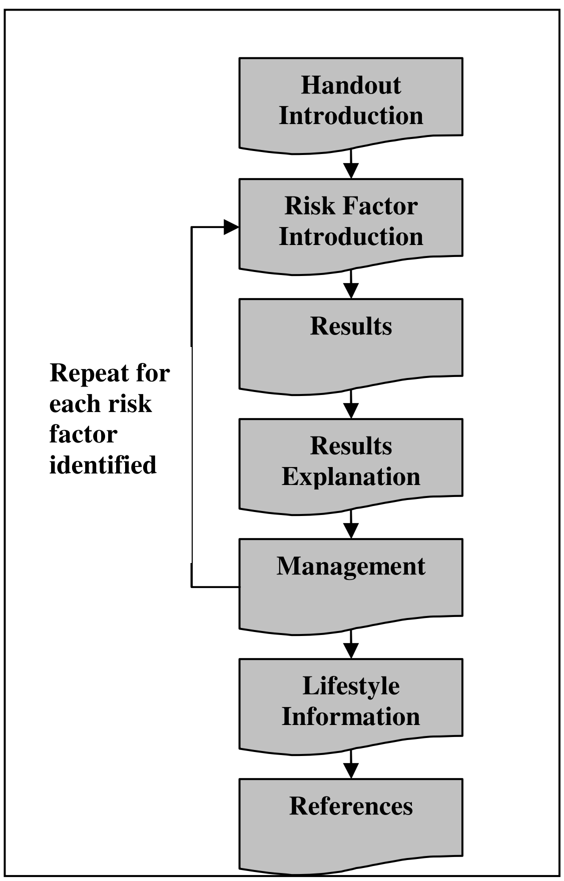

Figure 1. Document structure

In the risk factor section, each risk factor identified for a patient is discussed in turn as a separate topic. Each of these topics is subdivided into four smaller parts. An introduction section gives a brief overview of that particular risk factor. The results section provides that patient's own results, such as his or her LDL-C cholesterol level, and gives the target value for comparison. The explanation section interprets those results as normal or abnormal and the management section provides basic information on how to address the risk factor. All patients are provided additional detailed information about healthy diet and exercises habits in the lifestyle section. Finally, the references section describes where the information content was obtained.

For our application there were 63 different information packets in the message library. In order to personalize the packets as much as possible, each was stored as an XML document. By using XML tags such as <patient name> and <LDL-C>, we could easily insert personalized data into the pre-written text of the information packets. Once all packets were appropriately tagged, they were stored in the message library.

3.2.4 Algorithm. The fourth component needed for our computer-tailoring system was the algorithm, or the decision rules that direct which information packets the application should select and present based on a patient's data.. These rules are derived from the algorithm in the original CPG but are not identical to it. The CPG's algorithm uses evidence-based decision variables and their categorical values to direct recommendations for the physician. We had to create a series of rules that would, instead, permit the information tailoring system to select the appropriate information packets based on a patient's profile. As described above, we used a series of variables from the CPG that form a patient's profile. These variables and their accepted categorical values formed the "if" side of the rule statements. We subsequently created the "then" side of each statement, indicating to the application which information packet it should select given the data entered.

To facilitate this task, a decision tree was created. The decision variables formed the nodes of the tree and the information packets formed the leaves. The categories for the branches at each node were obtained from the original CPG. For example, according to the guideline the target level for LDL-C in a high risk person is less than 2.5 mmol/L. This established value was used by our tailoring application to determine whether a patient receives information about that risk factor. If a person's level is less than 2.5 mmol/L a packet is provided that congratulates him or her. If the level is greater than 2.5 mmol/L, information packets are provided that discuss the risk factor in more detail. In this way, the decisions for selecting the patient information materials are derived from the evidence-based knowledge in the CPG. The rules were written as simple "if-then" statements. For example, the rule for creating an information document on central obesity for a woman would be:

"IF ((dv1=female) AND (dv2>65) AND (dv7>88))
THEN (packet 1.00 AND packet 2.06 AND…)".

This directs the application to select and package 4 individual information packets into one document, related to managing central obesity in a woman with a waist circumference greater than 88 cm who is a senior. The decision rules were marked up with <logic> tags and stored together as an XML document for our document selection application.

3.2.5 Channel. The fifth and final component of our tailoring system was the channel for disseminating the messages. For the physician reminders, the application produces messages displayed on the computer screen in the web browser. The message delivered is the text of the action statement or directive for a given conditional or imperative recommendation respectively. Each tailored document is rendered in HTML on screen. For the upcoming pilot study we chose to provide paper documents that are mailed to patients. This will allow us to include individuals in the study who do not have access to a computer or are uncomfortable using a computer.

4. Implementation

Implementation of the GEM guideline was done by processing it with a Cocoon-based application. This application consists of two main components: a GEM Implementation System and the Document Selection Application (fig. 2).

4.1 GEM implementation system

The GEM implementation system is used to make the guideline case specific by allowing clinical data to be entered and showing relevant results. It is accessed through a web browser and works by dynamically creating input forms for each recommendation in the guideline. This is done by applying an XSLT stylesheet onto the underlying GEM document. Depending on the data type for that decision variable, the data is entered in either a text box or a radio button. When a user submits the completed form another XSLT stylesheet is applied. This stylesheet takes the clinical data and passes it into a rules engine. In addition to this data, the second XSLT stylesheet also looks at the relevant <logic> tags in the GEM document and passes the logic statements to the rules engine. The rules engine then processes the data based on this logic to determine the appropriate reminder (action statement) to display in the browser. This is done for each recommendation in the guideline.

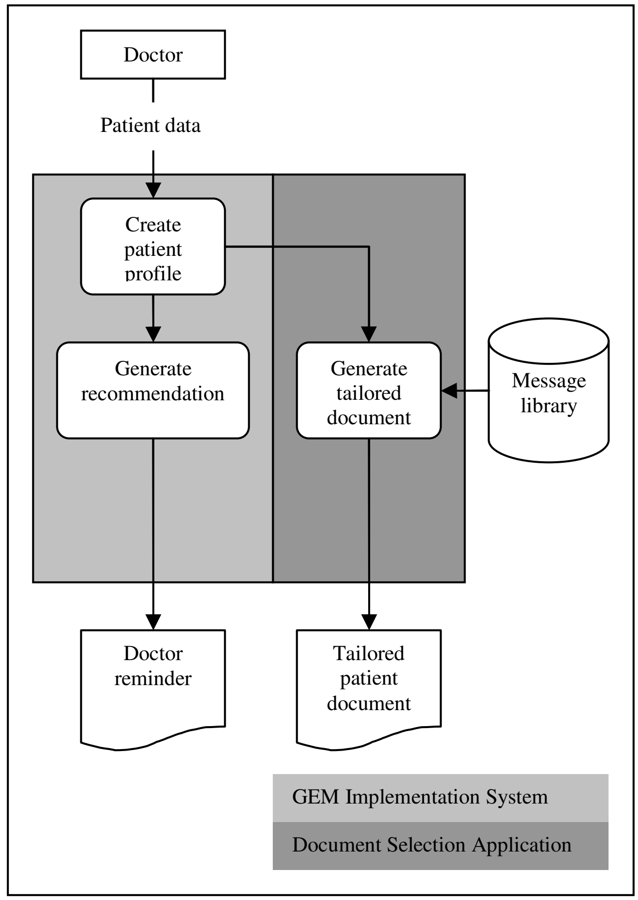

Figure 2. System design

4.2 Document selection application

When all the recommendations have been processed, the data collected and generated in the GEM implementation system (the patient's profile) are passed to our document selection application. We developed this application in XML for two reasons. First, if the information packets or the logic changes, a programmer would not necessarily be required to make the alterations to the code. If the application had been developed in Java, for example, a clinician would likely not be able to do this alone; however, most physicians are more likely to be able to work in XML. Second, we could use the rules engine described above to process the data and logic. In the document selection application, an XSLT stylesheet takes the data and the logic statements needed to select the documents and passes them to the rules engine. The rules engine uses this to determine the appropriate information packets to assemble from the message library and the personalized data to insert into those documents. The final result is a complete document tailored to each patient based on his or her own data.

5. Working example

A hypothetical patient can be used to illustrate the functionality of the system.

Mr. Smith is a 52 year old man with a total cholesterol level of 6.4 mmol/L, an HDL-C level of 1.2 mmol/L, an LDL-C level of 3.3 mmol/L and a triglyceride level of 1.8 mmol/L. His most recent blood pressure was 150/92 mmHg and he is being treated for hypertension. Mr. Smith is a smoker. He is not diabetic. He does not have a personal history of cardiovascular disease but his brother had a myocardial infarction at age 49. His abdominal circumference is 98 cm.

Mr. Smith's physician accesses the system through a web browser. When she does this, the GEM implementation system dynamically creates a series of input forms from the GEM-based CPG (fig.3). These forms direct her to enter the appropriate data about the patient for each of the recommendations in the guideline.

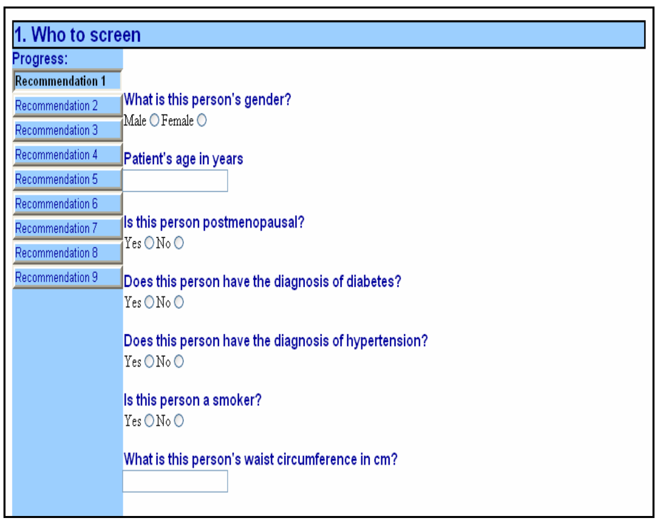

Figure 3. Data input form

For example, the second recommendation deals with calculating the patient's 10-year risk of coronary artery disease. The GEM implementation system displays a form asking for the values for each of the variables needed to complete the calculation. The physician submits this data about Mr. Smith and the system returns an action statement to the browser saying this man's 10-year risk is 12%. The physician works her way through each of the recommendations, entering data and receiving action statements, until they are done.

The clinical and demographic data the physician has been instructed to enter, along with certain values generated by the system (such as that for risk category), form Mr. Smith's personal profile. This collection of data is then passed to the document selection application where the rules engine compares it to the logic statements derived from the CPG and compiles the relevant information packets from the message library. In this example, Mr. Smith receives a personalized document consisting of 24 packets discussing containing information relevant to him

We have designed a pilot study to assess actual patient satisfaction with the materials that are generated for them. In this cross-over design trial using our system, tailored information materials will be generated for 25 patients based on their data supplied by their family physicians. Patients will be asked to review their own tailored materials as well as generic handouts that discuss their risk factors. They will be asked to complete a brief survey to elicit their satisfaction level with each of the materials (based on a number of variables) and to determine their preference for one format over the other.

6. Conclusions

Clinical practice guidelines are repositories of evidence-based medical knowledge. GEM was developed as a way to convert that knowledge into a format that can be processed by a computer. Once it can be processed, a guideline has an increased ability to provide decision support to physicians, and this type of reminder system has been shown to be effective in CPG implementation. The GEM-based guideline can also become something it was not intended to be: it can be the basis for generating tailored patient education materials. This type of patient-directed initiative is also a proven guideline implementation strategy and tailored education materials are more likely to be read and remembered, saved and discussed, and perceived as personally relevant [25]. By using a CPG as the knowledge base for this type of system, tailoring of the information materials is guided by the best clinical evidence available.

We have demonstrated one method of accomplishing this using a GEM-based guideline. Future work will use the information gained from this project and evaluate patient outcomes, including cholesterol levels and coronary events, in patients exposed to the tailored education materials generated by our system. Future research will also attempt to establish a method to effectively and efficiently integrate this system into physician's workflow patterns. According to James [31], in order to achieve wide acceptance by clinicians, CDS systems need to present the right information, in the right format, at the right time, without requiring special effort. As it was designed for this project, our system requires the physician to manually input all of the required data. Work is needed to mesh our system with an electronic medical record so that patient data can be drawn directly from it.

7. References

1. Haynes RB. What kind of evidence is it that Evidence-Based Medicine advocates want health care providers and consumers to pay attention to? BMC Health Services Research. 2002 Mar 6;2:1-7.
2. Jadad RA, Haynes RB. The Cochrane Collaboration—Advances and Challenges in Improving Evidence-based Decision Making. Med Decis Making. 1998;18:2-9.
3. Ray-Coquard I, Philip T, Lehmann M, Fervers B, Farsi F, Chauvin F. Impact of a Clinical Guidelines Program for Breast and Colon Cancer in a French Cancer Center. JAMA. 1997;278:1591-1595.
4. Barton S. Using clinical evidence. BMJ. 2001 Mar 3;322:503-4.
5. Eccles MP, Grimshaw JM. Selecting, presenting and delivering clinical guidelines: are there any "magic bullets"? MJA. 2004;180:S52-S54.
6. Shiffman RN, Karras BT, Agrawal A, Chen R, Marenco L, Nath S. GEM: A Proposal for a More Comprehensive Guideline Document Model Using XML. J Am Med Inform Assoc. 2000;7:488-498.
7. Grol R, Wensing M. What drives change? Barriers to and incentives for achieving evidence-based practice. MJA. 2004 Mar 15;180(supplement):S57-S60.
8. Vale MJ, Jelinek MV, Best JD, Santamaria JD. Coaching patients with coronary heart disease to achieve the target cholesterol: A method to bridge the gap between evidence-based medicine and the "real world"—randomized controlled trial. Journal of Clinical Epidemiology. 2002;55:245-252.
9. Lobach DF, Hammond E. Computerized Decision Support Based on a Clinical Practice Guideline Improves Compliance with Care Standards. Am J Med. 1997;102:89-98.
10. Shiffman RN, Karras BT, Agrawal A, Chen R, Marenco L, Nath S. GEM: A Proposal for a More Comprehensive Guideline Document Model Using XML. J Am Med Inform Assoc. 2000;7:488-498.
11. Grol R, Grimshaw J. From best evidence to best practice: effective implementation of change in patients' care. The Lancet. 2003 Oct 11;362:1225-1230.
12. Lemelin J, Hogg W, Baskerville N. Evidence to action: a tailored multifaceted approach to changing family physician practice patterns and improving preventive care. CMAJ. 2001 Mar 20;164(6):757-763.
13. Michie S, Johnston M. Changing clinical behaviour by making guidelines specific. BMJ. 2004 Feb 7;328:343-45.
14. Grimshaw J, Thomas R, MacLennan G, Fraser C, Ramsay C, Vale L, et al. Effectiveness and efficiency of guideline dissemination and implementation strategies. Health Technology Assessment. 2004;8(6):1-102.
15. Eccles M, McColl E, Steen N, Rousseau N, Grimshaw J, Parkin D et al. Effect of computerized evidence based guidelines on management of asthma and angina in adults in primary care: cluster randomized controlled trial. BMJ. 2002 Oct 26, pp. 325:1-7.
16. Hunt DL, Haynes RB, Hanna SE, Smith K. Effects of computer-based clinical decision support systems on physician performance and patient outcomes. JAMA. 1998;280:1339-1346.
17. Peleg M, Tu S, et al. Comparing computer-interpretable guideline models: A case-study approach. J Am Med Inform Assoc. 2003; 10: 52-68.
18. Agrawal A, Shiffman RN. Using GEM-encoded guidelines to generate medical logic models. Proc AMIA Symp. 2001;:7-11.
19. Shiffman RN, Agrawal A, Deshpande AM, Gershkovich P. An approach to guideline implementation with GEM. Medinfo. 2001;10(Pt1):271-5.
20. Guideline Elements Model. 2004. Available at: http://ycmi.med.yale.edu/GEM/. Accessed September 12, 2004.
21. Lorig K, Prohaska TR, Harris M, Gonzalez VM, Doak C, Doak L et al. Patient education: a practical approach. 3rd ed. Thousand Oaks, CA: Sage Publications; 2001. p. 21-55.
22. Tang PC, Newcomb C. Informing patients: a guide for providing patient health information. JAMIA. 1998 Nov/Dec;5(6):563-70.
23. Hobbs FDR, Erhardt L. Acceptance of guideline recommendations and perceived implementation of coronary heart disease prevention among primary care physicians in five European countries: the Reassessing European Attitudes about Cardiovascular Treatment (REACT) survey. Family Practice. 2002;19(6):596-604.
24. Mead VP, Rhyne RL, Wiese WH, Lambert L, Skipper B. Impact of environmental patient education on preventive medicine practices. The Journal of Family Practice. 1995;40(4):363-9.
25. Kreuter MW, Wray RJ. Tailored and targeted health communication: strategies for enhanicing information relevance. Am J Health Behav. 2003;27(Suppl 3):S227-S232.
26. De Vries H, Brug J. Computer-tailored interventions motivating people to adopt helath promoting behaviours: introduction to a new approach. Patient Education and Counseling. 1999;36:99-105.
27. Brug J, Oenema A, Campbell M. Past, present, and futre of computer-tailored nutrition education. Am J Clin Nutr. 2003;77(Suppl):1028S-34S.
28. Genest J, Frohlich JJ, Fodor JG, McPherson R, et al. Recommendations for the Management of Dyslipidemia and the Prevention of Cardiovascular Disease: 2003 update. CMAJ. 2003; 169 (9): online 1-10.
29. Karras BT, Nath SD, Shiffman RN. A preliminary evaluation of guideline content mark-up using GEM—an XML Guideline Elements Model. Proc AMIA Symp. 2000;:413-7.
30. Georg G, Seroussi B, Bouaud J. Interpretative framework to chronic disease management to guide textual guideline GEM-encoding. Stud Health Technol Inform. 2003;95:531-6.
31. James BC. Making it easy to do right. N Engl J Med. 2001;345:991-2.