A Flexible Approach to Guideline Modeling

Samson W. Tu, M.S., Mark A. Musen, M.D. Ph.D.
Stanford Medical Informatics, Stanford University School of Medicine
Stanford, CA 94305-5479, USA

We describe a task-oriented approach to guideline modeling that we have been developing in the EON project. We argue that guidelines seek to change behaviors by making statements involving some or all of the following tasks: (1) setting of goals or constraints, (2) making decisions among alternatives, (3) sequencing and synchronization of actions, and (4) interpreting data. Statements about these tasks make assumptions about models of time and of data abstractions, and about degree of uncertainty, points of view, and exception handling. Because of this variability in guideline tasks and assumptions, monolithic models cannot be custom tailored to the requirements of different classes of guidelines. Instead, we have created a core model that defines a set of basic concepts and relations and that uses different submodels to account for differing knowledge requirements. We describe the conceptualization of the guideline domain that underlies our approach, discuss components of the core model and possible submodels, and give three examples of specialized guideline models to illustrate how task-specific guideline models can be specialized and assembled to better match modeling requirements of different guidelines.

INTRODUCTION

Professional organizations, government agencies, and health-care institutions have published a plethora of clinical guidelines. The variety and complexity of guidelines presents a problem for researchers wishing to model them for the purpose of providing decision support. In this paper, we analyze the dimensions along which guidelines may vary, and describe a task-oriented approach to guideline modeling that we have been developing in the EON project.[1] In our approach, we manage complexity and variability of guidelines by extending a core guideline model with submodels, resulting in classes of guideline models that are matched to the knowledge requirements of different guidelines. We give examples of three classes of guidelines that we have modeled, describe the conceptualization of the domain that underlies our approach, discuss components of the core model and possible submodels that can be used for different tasks, and illustrate how task-specific models can be applied to the three example guidelines.

THREE EXAMPLES

We illustrate the variety of guidelines by three examples. An influenza vaccination guideline makes statements about how the decision of whether to vaccinate should be made. A simplified influenza guideline might state:

> *Patients at high risk of developing serious complications as result of influenza infection should be vaccinated. Patients at high risk include patients older than 65 years, and anyone with chronic diseases of the heart, lung, or kidney.*

A guideline for managing asthma in adults is a multi-encounter guideline for chronic disease that often specifies tasks to be performed and management decisions to be made on each encounter. A guideline for managing adult asthma might state:

> *On each encounter, check compliance with medications, take asthma history, record peak flow, look at asthma diary, check inhaler technique, and assess asthma state. For patients taking short-acting β_2 agonists and low-dose steroid inhalers, if asthma is not under control, consider stepping up to a medium-dose of steroid or adding a long-acting β_2 agonist.*

Finally, a breast-cancer clinical-trial protocol describes complex alternative treatment regimens for different groups of patients. As a prescriptive protocol, it usually specifies data-collection and treatment actions that are sequenced over time. Part of a clinical-trial protocol states:

> *Group I patients will receive Adriamycin 60mg/m^2 IV every 21 days for 4 cycles, along with Cytoxan 600mg/m^2 IV every 21 days for 4 cycles. Patients who are estrogen-receptor positive will receive tamoxifen PO for 5 years. Delay administration of Adriamycin and Cytoxan if there exists > grade 1 granulocytopenia on day 1.*

In the rest of the paper, we use these examples to illustrate our modeling framework.

DIMENSIONS OF A GUIDELINE MODEL

To build guideline-based applications, a developer has to create modeling concepts appropriate for the requirements of those applications. In the EON project, we are primarily interested in using guidelines to provide patient-specific decision support. To that end, we propose six dimensions along which modeling requirements of a guideline can be analyzed.

1. Provider behaviors that a guideline influences. We classify behaviors that a guideline tries to influence as (1) setting goals or constraints (e.g., "minimize HIV viral load"), (2) choosing an alternative among competing options (e.g., step up the dose of inhaled steroid versus

adding another asthma medication), (3) sequencing a set of actions (give 4 cycles of Adriamycin and Cytoxan, followed by 4 cycles of taxol), and (4) interpreting data (e.g., if the white-blood-cell count is between 2.0 and 2.9 thousand/mm^3, then the patient is experiencing Grade-2 granulocytopenia). A guideline may specify more than one kind of behavior. For example, decisions involve making choices that result in actions that need to be carried out as part of clinical workflow.

2. Temporal dimensions of actions and data. Actions modeled in a guideline may involve one-shot decision making (e.g., deciding whether to vaccinate against influenza), episodic interventions (e.g., outpatient encounter-based interactions between clinicians and asthma patients), continuous monitoring (e.g., adjusting ventilator settings in an ICU). In simple one-shot decision-making, the data needed may involve only current data. Alternatively, a guideline may use only time-stamped laboratory data, or it may require reasoning about temporal intervals.

3. Abstractions. Guidelines may require abstraction of data (e.g., interpreting white-blood-cell count as granulocytopenia), or abstractions of actions (e.g., prescribe a course of low-dose steroids without specifying the particular steroid medication).

4. Degrees of uncertainty. Guidelines differ regarding their data requirements and the degree of certainty that they place in their recommendations. Some, like clinical-trial protocols, are very explicit concerning the data that must be collected and the actions that must be carried out. Other guidelines merely are suggestive.

5. Point of view. Guidelines can be written for each class of participant in the health-care process: physicians, nurses, patients, or hospital administrators. Alternatively, a guideline can be written from the point of view of implementing a process of care.

6. Normal case and exceptions. As a way to manage complexity, one alternative for authors is to describe guidelines in terms of "normal" cases into which most target patients can be classified and "exceptional" cases that require special consideration. For example, some percentage of patients receiving a course of chemotherapy may develop several types of unusually severe reactions that require dose attenuation or hospitalization. In clinical protocols, details of how to manage such cases are often specified in special sections separate from the main protocol.

A developer of guideline-based application programs can analyze guidelines along these six dimensions to assess the modeling requirements. For example, a breast-cancer clinical trial (1) involves data interpretation and sequencing of tasks, (2) requires time-stamped patient data and interval-based activities (e.g., a course of chemotherapy), (3) abstracts laboratory tests results into toxicity intervals, and (4) describes management actions in the case of exceptional drug toxicities. An influenza vaccination guideline like the one described above, on the other hand, involves a one-shot decision requiring only knowledge of a patient's age and information about possible chronic conditions. A model that includes exception handling and complex action sequencing constructs therefore may not be appropriate for guidelines such as influenza vaccination. Thus, we see the need for an extensible framework for modeling clinical guidelines that allows addition of submodels to a core model. For such an approach to work, the core model and submodels must be coordinated through a common conceptualization of the domain.

CONCEPTUALIZATION OF THE DOMAIN

We conceptualize the guideline domain as consisting of multiple agents—healthcare providers, patients, and decision-support systems—interacting with each other at different points along a temporal continuum (Figure 1). We call these points of interaction *encounters*, although these points of encounter may simply be times when a monitoring system identifies the arrival of new data. At these encounters, observations about the patient are recorded, decisions are made, and actions are carried out. Actions are instantaneous acts—such as displaying a message to the user, making a referral, or starting a drug regimen—that lead to changes in the state of the world. Some actions change the state of on-going *activities*, such as the administration of a drug over time. Health-care providers and patients can take actions outside encounters (e.g., a patient may stop taking a drug).

GUIDELINE MODEL

Based on the conceptual model, we have developed a guideline model that can be partitioned into modular submodels and that allows alternative specializations of its modeling primitives. Corresponding to the basic

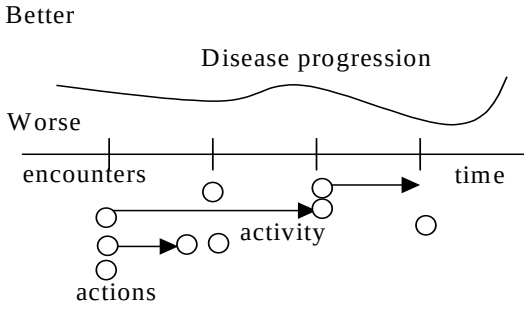

Figure 1. Schematic model of multi-encounter patient management. Decisions are made during encounters between healthcare providers and patient. Actions (represented as ovals), such as writing a prescription or having blood drawn for a lab test, are carried out during encounters. Some actions can start *activities* (represented as arrows) that extend over time.

distinction of time point and time line, we define two classes of guidelines: c*onsultation guidelines* that specify actions and decisions whose consequences are not being tracked over time, and *management guidelines* that model decisions and actions that lead to dependent changes in patient states over time. We define clinical algorithms that can use different types of decision-making mechanisms and alternative modeling primitives for expressing sequencing and repetition of actions. We distinguish between instantaneous *actions* and persistent *activities*, and we allow a guideline to be decomposed into statements about *normal cases* and *exceptions*. By taking combinations of these modeling primitives, we construct models appropriate for different classes of guidelines.

Clinical algorithm We model the decision-making and action-sequencing aspects of a guideline as parts of *clinical algorithms*. A clinical algorithm consists of a set of scenarios, action steps, decisions, and branching, synchronization, and repetition control nodes that are related by a *followed_by* relation. Decisions and action steps are *steps* in the clinical algorithm. Each step can have a goal associated with it.

A clinical algorithm can be specified for either consultation guidelines or management guidelines. When management recommendations of a guideline span multiple encounters, actions described in the guideline can be partitioned into those that have implications over time and those that constitute best-practice consultation actions recommended that encounter. The latter steps can be modeled in separate consultation guidelines that are indexed by the scenarios at which they apply.

Scenario A scenario is a partial characterization of the state of a patient. For example, the asthma medications that the patient is taking, such as low-dose inhaled steroids, may define a patient scenario. A scenario has an informal textual description and eligibility conditions that specify the necessary conditions for a patient to be in this scenario. Scenarios allow a clinician to synchronize the management of a patient to situations handled by a guideline. A scenario is followed by a decision or an action step.

Decision A decision represents a choice from a set of competing alternatives. A guideline model should be able to support different ways of describing alternatives and the corresponding selection mechanism. In the current model, we have two subclasses of decisions: decisions resolved by if-then-else conditions and decisions that require making a heuristic choice from a set of pre-enumerated alternatives. Making a choice among the alternatives is aided by preferences as determined by *rule-in* and *rule-out* conditions. If a *rule-out* condition evaluates to *true*, then an alternative is rejected. If the rule-out condition does not apply and a rule-in condition evaluates to *true*, then the alternative is marked as preferred. If neither evaluates to *true*, then the preference for the choice can be determined by a default preference associated with the alternatives.

Action sequencing and synchronization An action step is either a set of action specifications or a subguideline that has been sequenced in the clinical algorithm. The *followed_by* relation among decisions and action steps describes the sequencing of tasks in a clinical algorithm.

An action step describes a set of (instantaneous) actions that should be carried out. Sometimes it is necessary to describe decisions or sets of actions that should be carried out in separate but concurrent threads of execution (e.g., by different health-care providers). The model provides optional branching and synchronization primitives to allow specification of such parallel actions.

To model repetitive actions such as giving a fixed number of courses of chemotherapy cycles, we introduce special modeling constructs that specify the number of times or the frequency at which a set of actions should be carried. Like branching and synchronization, these modeling constructions are optional.

Goal In many guidelines, goals are not explicitly specified. Conversely, some guidelines may specify only the goals to achieve, rather than well-defined decisions to make or actions to perform. We associate goals (represented as Boolean criteria) with steps (i.e., with decisions and action steps), in the clinical algorithm.

Activity and Action Specifications Activities are processes that take place over time. They have states that can change from time to time. Such state changes can be the result of actions specified in a guideline. Some activities are abstract processes that are realized by more concrete processes.

The set of activities specified in the model is extensible. Currently we only model prescribable items that authorize the administration of drugs and related items. A drug regime models an abstract class of prescribable item (e.g., the class of possible calcium blocker prescription). The states of a drug regime are characterized by a set of attributes, such as dose level (e.g., low, high, medium) and frequency (e.g., twice a day).

Guidelines suggest instantaneous actions, including those that cause transitions among states of activities. An action can change the state of an activity by starting it, stopping it, or changing one of its attribute values. Other actions that can be specified in a guideline include acts such as collecting patient data, printing patient information leaflet, referring to other clinicians, or scheduling a follow-up appointment.

Data interpretation Data interpretation can take several forms in medicine. The EON guideline model currently supports three classes of data interpretation: (1) classification-based abstraction that recognize abstract concepts from concrete entities (e.g., zidovudine is a kind of anti-retroviral drug), (2) definition of terms based on certain criteria (e.g., if a person is older than 65 years then the person has high risk of influenza infection), and (3) knowledge-based temporal abstraction that creates interval-based abstractions from time-stamped data.[2] Each class of data interpretation is implemented in terms of specific classes of queries and criteria in the guideline model. Evaluation of these criteria causes appropriate interpreters to be invoked. Thus, for example, evaluating a temporal criterion that involves abstract concept may invoke a temporal mediator that creates these interval-based abstractions.

Exceptions For the clinical circumstances about which a guideline may have something to say, it may be impractical to express everything in clinical algorithms. A guideline author may want to partition the guideline into normal situations that cover usual cases and exceptional situations that rarely occur. In our model, we have defined two classes of exceptions: (1) exceptions that are *repairable* (i.e., those that lead a patient back to a scenario covered by the guideline), and (2) exceptions that are not repairable, and the patient is managed outside the guideline.

Just as we associate a consultation guideline with a scenario, we associate a set of exception handlers with a scenario. Before the execution engine presents choices or actions that follow a scenario, it checks to see if any current activity is generating exceptions, and invokes the appropriate exception handler. If exception handling activities do not result in a separate scenario, then these activities are considered to be part of the activities in the scenario; otherwise the exception-handling activities replace those of the current scenario.

APPLICATION OF GUIDELINE MODEL

For different classes of clinical guidelines, we derive different guideline models from the basic core model and various specializations of the model. All specialized guideline models share the distinction between actions and activities and include common models of time and of patient data, but they differ in the decision-making, action sequencing, and data interpretation primitives used in those models.

For one-shot guidelines such as the influenza-vaccination guideline, we defined a guideline model consisting of a consultation guideline that uses *if-then-else* condition for decision making, instantaneous actions that do not lead to activities with states, and state definitions (i.e., high-risk person) that use only current data. We use the Protégé knowledge modeling tool developed at our

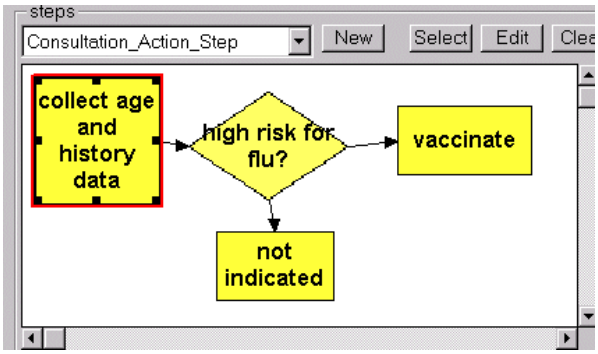

Figure 2 Influenza-vaccination guideline. "High risk for influenza" is defined as "age > 65 or presence of chronic heart, pulmonary, or x problems" The guideline uses one *if-then-else* constructs to model the decision to vaccinate

laboratory[3] to generate a guideline editor that allow such simple guidelines to be entered quickly (Figure 2).

In the case of asthma guideline, we define a guideline model that can be used to model chronic disease management in the manner of PRODIGY Phase 3 model.[4] Recommended data-collection actions at each encounter are modeled as consultation guidelines that are associated with scenarios in the management guideline. The management guideline defines scenarios in terms of the asthma medications that a patient is taking, and uses *if-then-else* and heuristic choice constructs to model guidance for decision-making. When a recommended action (such as starting a drug regime), is confirmed by a user, a corresponding activity is recorded as having started.

The model for clinical-trial protocols uses the full range of modeling constructs in the EON model. Statements such as "4 cycles of taxol" are represented as repetitions of actions involving the drug taxol as a prescribable item. The distinction between the clinical algorithm that a typical patient goes through (e.g., 4 AC cycles followed by 4 taxol cycles, with concurrent tamoxifen), and the section describing how the treatment should be altered when a patient experiences unusual drug toxicities (e.g., "if the patient experiences Grade-2 granulocytopenia, then delay administration of the chemotherapy AC") are modeled as management algorithm (Figure 3) and exceptions. Granulocytopenia is defined as an abstraction based on white-blood cell count.

DISCUSSION

Over the years, a variety of approaches for formal representation of clinical guidelines have been proposed and implemented. Computable representations of guidelines often assume particular formalisms, such as decision tables[5] or rules.[6] Even though a computational formalism may be natural for encoding a particular class of clinical guidelines, tying guideline representation so

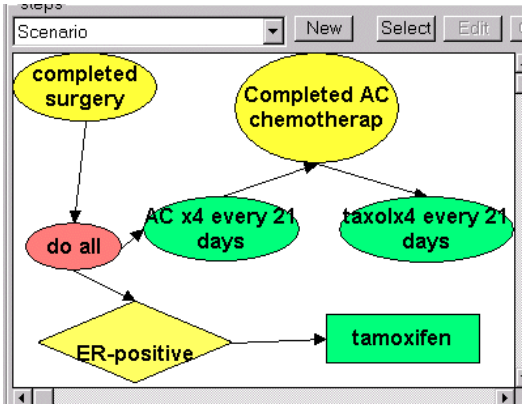

Figure 3. Part of management algorithm for a breast-cancer protocol. The arrows indicate a *followed-by* relation among decisions and actions specified in the protocol.

closely to computational formalism limits the generality of the model.

Several model-based approaches for representing clinical guidelines have been proposed. Guideline Interchange Format (GLIF) views a guideline as a plan of action with branching logic.[7] However, the GLIF model is very much under-specified. For example, it has no formal language for defining a Boolean criterion. The PRESTIGE project,[8] on the other hand, has developed a comprehensive and detailed model. However, the dynamics of the guideline representation depend on transition networks that relate states of various protocol uses and acts. Authoring a guideline means expressing much of the guideline content in terms of nodes, transitions, and transition criteria of the network. The complexity of the authoring process has proved to be a bottleneck.

We came to the conclusion that a monolithic model is unlikely to contain the exact distinctions required for modeling different classes of clinical guideline. A scalable and usable guideline model must have mechanisms for dealing with the variety and complexity of guidelines. In the EON project, we have tried to manage the complexity by defining a common conceptual model and developing a set of modeling primitives that can be mixed and matched to construct specialized models appropriate for different classes of guidelines. We defined alternative decision-making and action-sequencing constructs; we allowed a guideline to be factored into statements concerning normal cases and those concerning exceptional cases; we provided alternative abstraction mechanisms, and we distinguish between actions that pertains to a point in time and actions that have consequences for the future. For simpler guidelines, such as influenza vaccination guideline, we created a simpler guideline model. For more complex classes of guidelines, such management of chronic diseases and breast-cancer clinical-trial protocols, we created more complex guideline models to match the more demanding knowledge requirement.

Our work also speaks to the current effort to propose standards for guideline sharing. The definition of interchange format requires a prior consensus on a conceptual model that takes into account the variety of tasks and circumstances addressed by different guidelines. Our analysis of the dimensions of guidelines and the set of modeling constructs we have developed show why achieving such a consensus is so difficult.

ACKNOWLEDGMENTS

This work has been supported in part by grant LM05708 from National Library of Medicine and DARPA contract #N66001-94-D6052. We are grateful to Doug Fridsma, Peter Johnson, and Yuval Shahar for discussions related to the automation of guideline-based care. Samson Tu thanks Sowerby Centre for Health Informatics in Newcastle upon Tyne for hosting a productive summer sabbatical. A gift from NHS Executive to Stanford Medical Informatics made the sabbatical possible.

REFERENCES

1. Musen M, Tu S, Das A, Shahar Y. EON: A Component-Based Approach to Automation of Protocol-Directed Therapy. JAMIA 1996;3:367-388.

2. Shahar Y. A Framework for Knowledge-Based Temporal Abstraction. Artificial Intelligence 1997;90(1-2):79-133.

3. Musen MA, Gennari JH, Eriksson H, Tu SW, Puerta AR. PROTEGE II: Computer Support For Development Of Intelligent Systems From Libraries Of Components. MEDINFO 95, The Eighth World Congress on Medical Informatics. Vancouver, B.C., Canada, 1995:766-770.

4. Johnson P, Tu SW, Booth N, Sugden B, Purves I. A Guideline Model for Chronic Disease Management in Primary Care. Report SMI-1999-790, Stanford Medical Informatics, Stanford University School of Medicine, 1999.

5. Shiffman RN. Representation of Clinical Practice Guidelines in Conventional and Augmented Decision Tables. JAMIA 1997;4(5):382-393.

6. Hripcsak G, Clayton PD, Pryor TA, Haug P, Wigertz OB, van der Lei J. The Arden Syntax for Medical Logic Modules. Proceedings of the 14th Annual Symposium on Computer Applications in Medical Care. Washington, D.C., 1990:200-204.

7. Ohno-Machado L, Gennari JH, Murphy S, et al. The GuideLine Interchange Format: A Model for Representing Guidelines. JAMIA 1998;5(4):357-372.

8. Gordon C, Veloso M. The PRESTIGE Project: Implementing Guidelines in Healthcare. Medical Informatics Europe '96: IOS Press, 1996:887-891.