

The Structure of Guideline Recommendations: A Synthesis

Samson W. Tu, M.S.¹, James Campbell, M.D.², Mark A. Musen, M.D. Ph.D.¹

¹Stanford Medical Informatics, Stanford University School of Medicine, Stanford, CA

²University of Nebraska Medical Center, Nebraska Health Systems, Omaha, NE

ABSTRACT

We propose that recommendations in a clinical guideline can be structured either as collections of decisions that are to be applied in specific situations or as processes that specify activities that take place over time. We formalize them as “recommendation sets” consisting of either Activity Graphs that represent guideline-directed processes or Decision Maps that represent atemporal recommendations or recommendations involving decisions made at one time point. We model guideline processes as specializations of workflow processes and provide possible computational models for decision maps. We evaluate the proposed formalism by showing how various guideline-modeling methodologies, including GLIF, EON, PRODIGY3, and Medical Logic Modules can be mapped into the proposed structures. The generality of the formalism makes it a candidate for standardizing the structure of recommendations for computer-interpretable guidelines.

INTRODUCTION

Over the last twenty years, many groups have developed formalisms for implementing clinical practice guidelines (CPGs) and protocols in systems that provide patient-specific decision support at points of care. Most systems have been developed in academic settings and have not gained widespread acceptance. Only Arden Syntax for Medical Logic Modules (MLMs)¹ has achieved the status of an industry standard. Yet MLMs are designed for encoding the knowledge necessary to make single decisions. Many see the development of a shared standard for encoding complex multi-step CPGs as the next step to achieve economy in guideline knowledge-base development and to enhance the prospect of guideline implementation by vendors of clinical information systems.²

In 2000, the Health Level 7 (HL7) Clinical Decision Support Technical Committee created a special interest group (SIG) devoted to clinical guidelines. Participants of the SIG recognized that, while the messaging standards being developed in HL7 provide the opportunity for decision-support systems to achieve data interoperability with clinical information systems, no single guideline formalism has emerged as the consensus standard format. Consequently, the SIG decided that it can more effectively work on

standardizing *components* of a computable guideline formalism, such as an expression language for clinical decision support and the overall flow of recommendations in a guideline.

Recommendation statements represent the fundamental assertions of guidelines. Prior studies have shown that there are great deal of commonalities among the ways that guideline modeling methodologies organize recommendations³ The initial requirements for a standardized way of representing the overall flow of recommendations include (1) an integrated description of decision making and activity specification, (2) expressive process description that allows for sequencing, repetition, and concurrency of decisions and activities, and (3) clear semantics for the representation constructs. In the attempt to satisfy these requirements, we reviewed the relevant literature and now propose a synthetic model that distinguishes between those *recommendation sets* that include process descriptions and those that primarily define recommendations for making decisions at a point in time. We base these formulations in terms of well-understood process and computational models and evaluated them by showing how the control structure of several guideline-modeling methodologies can be mapped into this recommendation-set formalism.

METHOD

We analyzed several diagramming conventions⁴ and the organization principles of several different guideline modeling methodologies.⁵⁻¹⁰ As a result of the study we formulated four distinct paradigms for semi-formal and formal organization of guideline recommendations: (1) flowchart for human understanding, (2) decision maps that enumerates patient scenarios in which medical decisions are made, (3) plans that describe partially ordered actions designed to achieve certain goals, and (4) workflows that formulate guideline logic in terms of processes made of activities in which organizational actors play specific roles.¹¹ Among these four paradigms, flowcharts for human understanding, which are often found in books and medical journals, cannot be made computer-interpretable directly, as they assume that trained clinicians, with the aid of textual annotations associated with flowcharts, can interpret the intent of the algorithms, adjust for possible arbitrary ordering of the steps, resolve ambiguities, and fill in implicit knowledge assumed by flowchart developers. Planning

and workflow representations share the characteristic that they both describe *processes* made up of activities that take place over time. Consequently we propose that recommendations in a guideline can be organized as *recommendation sets* consisting of either *Activity Graphs* that represent guideline-directed care processes or *Decision Maps* that represent recommendations involving decisions at a time point.

Recommendation Set

We define a recommendation set for a computable guideline as a collection of related recommendations that are applicable in one or more shared contexts and that are organized according to a computable formalism. A context is defined by a combination of a clinical setting (e.g. outpatient encounter in a general internal medicine clinic), the care provider to whom the recommendation is directed, and the relevant patient states (e.g. a hypertensive patient who has been prescribed anti-hypertensive agents). Within each context, a recommendation may describe the preferred choice in a management decision (e.g. whether to increase the dose of a drug or to add another anti-hypertensive agent), or it may recommend a series of actions be carried out (e.g. perform history and physical before ordering certain tests).

Activity Graph

An Activity Graph describes the relationships among different activities in terms of a process model. Several process models have been proposed as standards in the literature. We adopted the Workflow Management Coalition's Workflow Process Definition Language (WPDL)¹² as the basis for the semantics of an Activity Graph. We do so because WPDL's object-oriented feature, the close correspondence between WPDL and the process specification of existing guideline models, prior experience with WPDL in the guideline-modeling community,⁷ and the available literature and software for workflow management systems.

In WPDL, a process consists of a collection of activities each of which is a unit of work that is carried out by a combination of organizational resources and/or computer applications. Activities are related through precedence relations as defined by transition and transition conditions, and by hierarchical decomposition relations as defined by subflow relationships. Sets of activities may be carried out sequentially, concurrently, or in any order, as indicated by *split* and *join* constraints. If the *split* constraint of an activity is AND, then all transitions to subsequent activities need to be evaluated for possible activation. If the *split* constraint is XOR, then the outgoing transitions are evaluated in order, and the first transition that evaluates to *true* is taken. Similarly, if the *join* is AND, all active preceding activities must be

completed and the transition conditions evaluated to *true* before the current activity can start. If the *join* constraint is XOR, then the current activity is enabled as soon as the first preceding activity completes and the associated transition condition evaluates to *true*. Thus, the AND *split* constraint allows concurrent activities to be started following an activity, and the *join* constraint synchronizes multiple threads of execution. Dummy activities that have no associated work (called Route activities in WPDL) can be used solely for specifying join and split constraints among other activities. Other information, such as priority and automation mode (i.e., whether the activity can be started or stopped automatically) can be attributes of an activity. The model also allows extensions through addition of attributes that are necessary for particular applications.

We define an Activity Graph as a specialization of the workflow process model for specifying clinical and computational activities designed to implement guideline recommendations. Activity Graphs represent related sets of actions and decisions as recommended in a particular clinical and organizational context. Formally, an Activity Graph is a connected directed graph of nodes and arcs, where nodes represent activities and arcs represent transitions among activities. Each node inherits from the workflow process model attributes such as join and split constraints and automation modes. As a process, it should have a well-defined starting node. To adapt the generic workflow processes to model guideline-directed processes, we specialize the Route activity of WPDL into a Context node or a Route node, and a generic workflow activity into an Action node or a Decision node in an Activity Graph.

From the point of view of workflow, a Context node only has routing information. However, we use the Context node to specify relevant information about trigger states, clinical setting, enterprise resources, providers to whom the recommendations are directed, and the assumed patient's clinical and management states. Figure 1 shows a simple Activity Graph that depicts the top-level care process during a routine outpatient encounter in which a guideline for the management of hypertension is implemented. The process starts when a nurse accesses a patient's medical record in an outpatient clinic. That event triggers the execution of the hypertension guideline's Activity Graph by a decision-support system in the context of "Nurse Interview."

An Action node encapsulates a set of work items that should be performed either by a computer system or by a healthcare provider. In Figure 1, the History and Physical and Adjust Medication Action nodes within the "Physician Interview" context contains recommendations that can be carried out in parallel

through an AND split. Both must be completed before the context switches to that of patient check-out.

The Decision node is the third major class of nodes in an Activity Graph. A Decision node represents a decision point in an Activity Graph. We will discuss details of a Decision node in the next section.

Decision Map

A Decision Map is a collection of recommendations that do not need to be organized and executed as a process. One use of the Decision Map is the encoding of a collection of asynchronous alerts and reminders that are not organized as a connected process of activities. Alternatively, a Decision Map may be used as the decomposition of a high-level action such as “Adjust Medication” in Figure 1. The high-level action involves decisions made by a single provider in a specific clinical context at a single time. Figure 2 shows a portion of the top-level recommendation structure of the ATHENA hypertension management system,¹³ which organizes choices on use of anti-hypertensive medication as a collection of decision points differentiated by a patient’s clinical and management states. Formally, Decision Map consists of a collection of decision points, each of which consists of a Context node, a Decision node, and a set of Action nodes. As in an Activity Graph, a Context node is characterized by triggers, clinical setting, patient state, and current therapy. An Action node specifies actions to be performed either by the system or by a healthcare provider. A Decision node is the locus of decision knowledge organized according to a *decision model*. Examples of decision models are decision criteria as in the logic slot of an MLM¹, for-and-against argumentation structures as used in the PROforma methodology,⁹ or decision trees and Bayesian nets as used in GUIDE.⁷ The contract between a recommendation set and a Decision node is that, at the end of the decision-making process, commitment to one or more alternatives are made.

Computationally a Decision Map allows several possible specializations. A Decision Map can be a collection of event–condition–action rules where the Context nodes define triggering events, the Decision node the condition, and the Action node the action. Alternatively, if triggering is defined externally, and if a Decision Map has the constraint that only one choice can be made in a decision, and an action always results in a new context, then computationally the Decision Map becomes an augmented transition net (ATN). Within this paradigm, each time the Decision Map is invoked, the guideline execution system determines the current context, applies the knowledge in the decision model, and commits to one alternative among the choices.

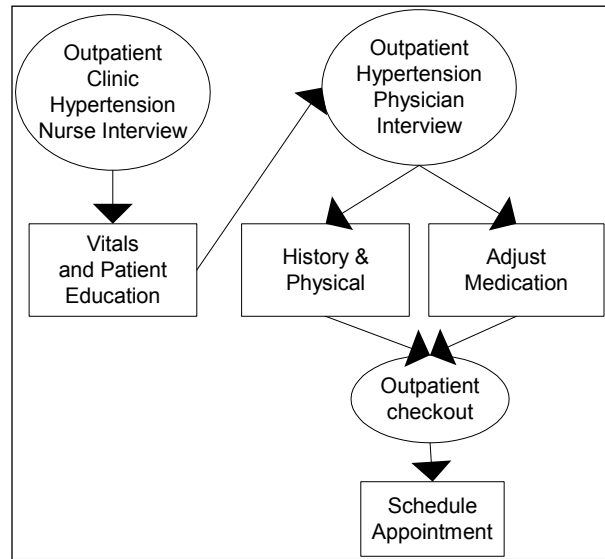


Figure 1. An Activity Graph depicting the contexts (ovals) in which guideline recommendations apply and actions (rectangles) that should be performed.

Sub-recommendation

In order to manage complexity in guideline recommendations, we need to support hierarchical nesting of recommendations. Each Context, Decision, and Action node in an Activity Graph or a Decision Map may be associated with another recommendation set that we call sub-recommendation. A sub-recommendation can be executed synchronously or asynchronously. A synchronous sub-recommendation means that the execution of the parent node is not completed until execution of the sub-recommendation is completed. Mandatory subtasks, for example, can be encoded as a synchronous sub-recommendation. Asynchronous sub-recommendation means that the parent node does not wait for the sub-recommendation to complete its execution. Instead, management of is forked off as an independent process.

As a matter of guideline encoding convention, a sub-recommendation in a Context node should recommend actions that are relevant in that context, regardless of any subsequent decisions or actions. A sub-recommendation in a Decision node should recommend actions that are helpful in making the decision (e.g. obtaining relevant information about patient state). Finally, a sub-recommendation in an Action node should refine the actions specified in the Action node.

EVALUATION

We evaluated this formulation of recommendation sets in terms of its ability to reproduce or clarify the structure of guideline recommendations in existing guideline modeling methodologies. So far, we have

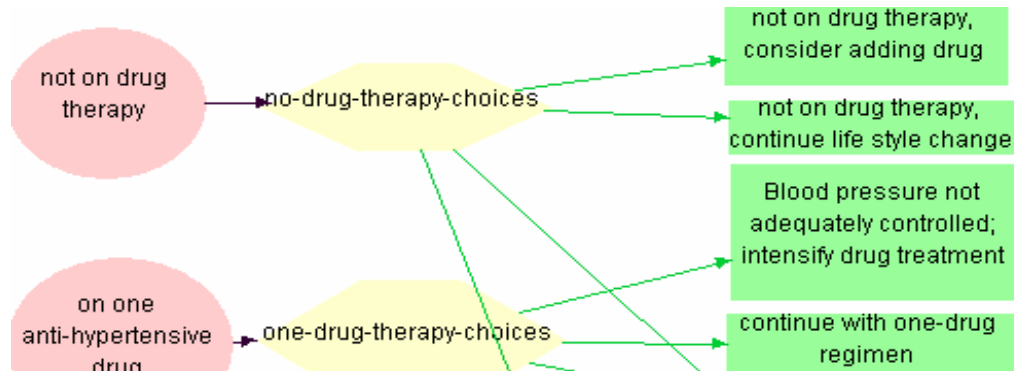


Figure 2. A Decision Map showing the Contexts (ovals) of decision points (hexagons) and some of the decision alternatives (green rectangles) at the two decision points.

established mappings from recommendation formalisms of Medical Logic Modules, PRODIGY3,⁸ EON,⁵ and GLIF¹⁰ into the recommendation set structure proposed in this paper. We have done partial mapping from the workflow constructs in the HL7 Reference Information Model.¹⁴ In doing so, we have uncovered discrepancies that need to be resolved for the HL7 RIM to provide the semantic basis for guideline recommendations.

The event–condition–action form of Decision Map maps directly into MLMs, where the event is the trigger, condition the logic slot, and action the action slot. A PRODIGY3 management-guideline diagram is essentially a Decision Map where the Context node and Decision node following it are conflated into a PRODIGY Scenario. An asynchronous Activity Graph (Consultation guideline) is associated with the Context node, and PRODIGY3 Action steps that expand into sub-guidelines correspond to Action nodes that have Decision Maps as sub-recommendations. The execution semantics of a PRODIGY3 Decision Map are strictly that of an augmented transition network.

EON and GLIF3 methodologies do not distinguish between Activity Graph and Decision Maps. EON’s Scenario step and GLIF’s Patient_State step are specializations of the Context node. Like PRODIGY3, EON’s Scenario node may have an associated Activity Graph. Subguidelines of GLIF’s Decision and Action steps correspond to Activity Graphs sub-recommendations. EON’s sub-guideline step maps to an Action node that has an associated Activity Graph. The Branching and Synchronization steps in the two methodologies correspond to Route nodes, where the Synchronization step in GLIF may expand into a sub-graph of Route nodes because the synchronization condition of a GLIF Synchronization step may have an AND/OR structure.

We had done some preliminary mapping between constructs in this proposal and the HL7 version 3 RIM.

A Context node is a complex object involving clinical settings (Place), participants (Roles), and patient clinical, administrative and management states (Acts). A Decision node, because its role in generating a commitment, can be seen as a derived Observation where the value of the Observation is the committed choice. Action and Route nodes are examples of HL7 Acts. HL7 RIM has a number of Act Relationships that are relevant to this proposal. The precondition Act Relationship relates an act with another act in the criterion mood. Thus, we can specify Boolean preconditions on decision alternatives. However, the RIM provides no role for complex decision models that may be necessary to encode guideline recommendations. The decomposition Act Relationship decomposes an Act into a collection of component acts that are numerically sequenced, where sub-acts having the same sequence number are to be carried out concurrently. This formulation does not allow for a sequence of acts to be carried out between branching and synchronization. If we encapsulate such a sequence with a dummy composite act, we see that the difference is only syntactic. Another difference is that, unlike the workflow process model where split and join restriction are placed on activities to govern the transitions into and out of the activity, the analogous split and join restrictions on the RIM are specified as attributes on the decomposition Act Relationship.

CONCLUSION

We have created a framework for organizing guideline recommendations into process-oriented Activity Graphs and decision-centric Decision Maps. For the former, we use the Workflow Management Coalition’s workflow process model as the underlying process definition. For the latter, we provided alternative computational methods. We evaluated the framework by mapping several guideline modeling formalisms to it and compared it to the workflow management aspect of the HL7 RIM. As we perform more detailed comparisons, we will discuss the results with

developers of the RIM and possibly make suggestions on how the two approaches may be reconciled. The generality of the formalism makes it a candidate for standardizing the structure of recommendations in computer-interpretable guidelines.

The primary benefit in making the distinction between process-oriented Activity Graph and decision-centric Decision Map is that we can support the distinction in terms of computational models that have well-understood semantics, which in turn allows us to show cleanly how the framework generalizes the those of other methodologies.

By itself, this framework of recommendation set does not constitute a guideline model in which clinical guidelines can be encoded. Here we are concerned with the *structure* of recommendations, i.e. high-level categorization of the components of a recommendation and how recommendations relate to each other. In the SAGE project, we use recommendation sets as the nucleus of a complete interoperable guideline modeling framework that includes alternative decision models, detailed action specifications, enterprise resources and roles, and standard data models and terminologies.¹⁵

ACKNOWLEDGEMENT

This work has been supported by the U.S. National Institute of Standards and Technology, Advanced Technology Program, Cooperative Agreement Number 70NANB1H3049, by grant LM05708 from National Library of Medicine, and by grant LM06594 with support from the Department of the Army, Agency for Healthcare Research and Quality, and the National Library of Medicine.

REFERENCES

1. Hripcsak, G, Clayton, PD, Pryor, TA, Haug, P, Wigertz, OB, Van der lei, J. The arden syntax for medical logic modules. Fourteenth Annual Symposium on Computer Applications in Medical Care; 1990 Washington, DC. 200-204.
2. Boxwala, A, Tu, SW, Zeng, Q, et al. Towards a representation format for sharable clinical guidelines. Journal of Biomedical Informatics 2001;34(3):157-169.
3. Peleg, M, Tu, SW, Bury, J, et al. Comparing computer-interpretable guideline models: A case-study approach. JAMIA 2003;10(1):52-68.
4. Society for Medical Decision Making(SMDM) Committee on Standardization of Clinical Algorithms (CSCA). Proposal for clinical algorithm standards. Medical Decision Making 1992;12(2):149-154.
5. Tu, SW, Musen, MA. A flexible approach to guideline modeling. Proceedings of AMIA Annual Symposium; 1999 Washington DC. 420-424.
6. Miksch, S, Shahar, Y, Johnson, P. Asbru: A task-specific, intention-based, and time-oriented language for representing skeletal plans. 7th Workshop on Knowledge Engineering: Methods & Languages (KEML-97); 1997 Milton Keynes, UK. 9-1-9-20.
7. Quaglini, S, Stefanelli, M, Cavallini, A, Micieli, G, Fassino, C, Mossa, C. Guideline-based careflow systems. Artificial Intelligence in Medicine 2000;5(22):5-22.
8. Johnson, PD, Tu, SW, Booth, N, Sugden, B, Purves, IN. Using scenarios in chronic disease management guidelines for primary care. Proceedings of AMIA Annual Symposium; 2000 Los Angeles, USA. 389-393.
9. Fox, J, Das, SK. Safe and sound. Cambridge: MIT Press; 2000.
10. Peleg, M, Boxwala, A, Ogunyemi, O, et al. Glif3: The evolution of a guideline representation format. Proceedings of AMIA Annual Symposium; 2000 Los Angeles. 645-649.
11. Tu, SW, Johnson, PD, Musen, MA. A typology for modeling processes in clinical guidelines and protocols. Stanford: Stanford Medical Informatics; 2002. Report No.: SMI-2002-0911.
12. Workflow Management Coalition. Interface 1: Process definition interchange, process model. Lighthouse Point, FL: Workflow Management Coalition; 1999 October 29. Report No.: WfMC TC-1010-P.
13. Goldstein, MK, Hoffman, BB, Coleman, RW, et al. Operationalizing clinical practice guidelines while taking account of changing evidence: Athena, an easily modifiable decision-support system for management of hypertension in primary care. Proceedings of AMIA Annual Symposium; 2000 Los Angeles, USA. 280-284.
14. Schadow, G, Russler, DC, McDonald, CJ. Conceptual alignment of electronic health record data with guideline and workflow knowledge. Int J Med Inf 2001;64:259-274.
15. Campbell, J, Tu, SW, Boyer, J, Mansfield, G. The sage guideline model: A knowledge representation framework for encoding interoperable clinical practice guidelines. Stanford: Stanford Medical Informatics; 2003 Report No.: SMI-2003-XXX.