Automated Support to Clinical Guidelines and Care Plans:
The Intention-Oriented View

Yuval Shahar M.D., Ph.D.,
Medical Informatics Research Center
Department of Information Systems Engineering
Ben Gurion University, Beer Sheva, Israel 84105

1. Clinical Guidelines: An Introduction

Clinical guidelines (or *Care Plans*) are a powerful method for standardization and uniform improvement of the quality of medical care. **Clinical guidelines** are a set of schematic plans, at varying levels of abstraction and detail, for management of patients who have a particular clinical condition (e.g., insulin-dependent diabetes). (C*linical protocols* are typically highly detailed guidelines, often used in areas such as oncology and experimental clinical trials.) The application of clinical guidelines by care providers typically involves collecting and interpreting considerable amounts of data over time, applying standard therapeutic or diagnostic plans in an episodic fashion, and revising those plans when necessary.

Clinical guidelines can be viewed as reusable *skeletal plans* that, when applied to a particular patient, need to be refined by a care provider over significant time periods, while often leaving considerable room for flexibility in the achievement of particular goals. Another view, which I will dwell upon in more length here, is that clinical guidelines are a set of *constraints* regarding the *process* of applying the guideline (i.e., care-provider actions) and its desired *outcomes* (i.e., patient states), to which I refer as process and outcome **intentions**. These constraints are mostly *temporal*, or at least have a significant temporal dimension, since most clinical guidelines concern the care of chronic patients, or at least specify a care plan to be applied over a significant period.

It is now universally agreed that conforming to state-of-the-art guidelines is the best way to improve the quality of medical care, a fact that had been rigorously demonstrated [Grimshaw and Russel, 1993], while reducing the escalating costs of medical care. Clinical guidelines are most useful at the point of care (typically, when the care provider has access to the patient’s record), such as at the time of order entry by the care provider.

2. The Need for Automated-Support to Clinical Guidelines

Most clinical guidelines are text-based and inaccessible to the physicians who most need them. Even when guidelines exist in electronic format, and even when that format is accessible online, physicians rarely have the time and means to decide which of the multiple guidelines best pertains to their patient, and, if so, exactly what does applying that guideline to the particular patient entail. Furthermore, recent health-care organizational and professional developments often reduce guideline accessibility, by creating a significant information overload on health care professionals. These professionals need to process more data then ever, in continuously shortening periods of time. Similar considerations apply to the task of assessing the quality of clinical-guideline application

To support the needs of health-care providers as well as administrators, and ensure continuous quality of care, more sophisticated information processing tools are needed. Due to limitations of state-of-the-art technologies, analyzing unstructured text-based guidelines is not feasible. Thus, there is an urgent need to facilitate guideline dissemination and application using machine-readable representations and automated computational methods.

Several of the major tasks involved in guideline-based care, which would benefit from automated support, include specification (authoring) and maintenance of clinical guidelines, retrieval of guidelines appropriate to each patient, runtime application of guidelines, and retrospective assessment of the quality of the application of the guidelines.

Supporting guideline-based care implies creation of a *dialog* between a care provider and an automated support system, each of which has its relative strengths. For example, physicians have better access to certain types of patient-specific clinical information (such as their odor, skin appearance, and mental state) and to general medical and commonsense knowledge. Automated systems have better and more accurate access to guideline specifications and detect more easily pre-specified complex temporal patterns in the patient’s data. Thus, the key word in supporting guideline-based care is *synergy*.

3. The State of The Art in Automation of Clinical Guidelines

Several approaches to the support of guideline-based care permit hypertext browsing of guidelines via the World Wide Web [Barnes and Barnett, 1995] but do not directly use the patient’s electronic medical record, and do not attempt to reduce the load on physicians by obviating the need for actually reading the guideline and customizing it to the patient’s personal clinical history and current state.

Several simplified approaches to the task of supporting guideline-based care, which *do* use the patient’s data, encode guidelines as elementary state-transition tables or as situation-action rules dependent on the electronic medical record [Sherman et. al., 1995]. An established (ASTM) medical-knowledge representation standard, the Arden Syntax (Hripcsak et al., 1994), represents medical knowledge as independent units called *Medical Logical Modules* (*MLMs*), and, quite importantly, separates the general medical logic (encoded in the Arden syntax) from the institution-specific component (encoded in the query language and terms of the local database). Rule-based approaches, how, do have several disadvantages: (1) They typically do not include an intuitive, explicit representation of the guideline’s overall clinical logic; (2) they have no semantic distinctions regarding the different types of clinical knowledge represented; (3) they lack the ability to easily represent and reuse guidelines and guideline components; (4) they cannot represent, use, and reuse higher, meta-level problem-solving knowledge; and (5) they do not support application of guidelines over extended periods of time, as is necessary to support the care of chronic patients guideline-based care over extended periods in automated fashion.

During the past 20 years, there have been several efforts to support complex guideline-based care over time in automated fashion as explicit, well-defined plans. Examples of architectures and representation languages include ONCOCIN [Tu et. al., 1989], T-HELPER [Musen et. al., 1992], DILEMMA [Herbert et. al, 1995], EON [Musen et. al., 1996], Asgaard [Shahar et al., 1998], PRO*forma* [Fox et al., 1998], the guideline interchange format (GLIF) [Ohno-Machado et al., 1998; Peleg et al., 2000], the European PRESTIGE project, and the British Prodigy project [Johnson et al., 2000].

Most of the approaches for supporting guideline-based care, and in particular the more disciplined ones, which are based on a planning paradigm, can be described as being *prescriptive* in nature, specifying *what* actions need to be performed and *how*. However, several systems, such as Miller's VT-Attending system, have used a *critiquing* approach, in which the physician suggests a specific therapy plan and gets feedback from the program. The Asgaard project uses both a prescriptive methodology for specification of prescribed interventions and a critiquing approach for retrospective quality assessment.

Several recent approaches to support guideline use at the crucial point of care, at which the care provider is most susceptible to advice, enable a Web-based connection from an electronic patient record to an HTML-based text guideline to which certain special annotations were added, enabling some form of execution of the guideline. A good example is the *ActiveGuidelines* model [Tang and Young, 2000], which is embedded in a well-known commercial electronic medical record system, but can be generalized to other electronic medical record systems, as long as these systems use a specialized *ActiveGuidelines Interpreter* that can interpreter hidden tags in the HTML guideline's representation, converting these to specific order-entry rules that link to the local electronic medical record. Although such approaches do have the advantage of simplicity and are quite useful for certain well-defined contexts, they usually have no standardized, sharable, machine-readable representation of guidelines that can support multiple tasks such as automated eligibility determination, customized application, and retrospective quality assurance, and are currently not intended for representation of complex care plans over time.

A recent framework, the *Guideline Elements Model* (*GEM*), enables structuring of a text document containing a clinical guideline as an *extensible markup language* (*XML*) document, a useful specification for sharing documents over the Internet, using a well-defined XML schema [Karras et al., 2000; Shiffman et al., in press]. GEM, however, is an application running on a stand-alone computer, and does not support any computational tools that can interpret the resulting semi-structured text, since it does not include a formal language that provides a clear computational model.

4. The Importance of Representing Guideline Intentions

Automating guideline-based care requires the use of an underlying richly expressive, machine-readable formal language, specific to that task, which enables specification of (1) multiple types of clinical actions over time (e.g., sequential, parallel, periodic) and associated temporal and other constraints (e.g., administration between 8:30 to 10:00 A.M.), and (2) the intermediate and overall clinical-processes and patient-outcome goals of the therapy plan, namely, the *process* and *outcome intentions* of the guideline. These intentions are temporal-pattern constraints (e.g., a process intention to administer regular insulin twice a day; an outcome intention to maintain fasting blood glucose within a certain range over at least 5 days a week) that have individual *weights* signifying their relative importance. Such knowledge is necessary to determine whether a care provider who has modified some portion of the guideline is still following most of the guideline, or, at least, its spirit, and to what extent. Such a provider might be applying the guideline, albeit in modified fashion, as is often the case. Intentions are also crucial for supporting a provider who needs to modify the guideline due to local constraints, but needs to know what was the design rational of the guideline to decide which modification is legitimate (thus avoiding undue alarm) and which would significantly harm either the intended process or intended outcome (thus insuring the integrity of the application and the safety of the patient). Changing the type of an anti-hypertension drug could be quite reasonable, or could be in conflict with either the process intention (as would be the case for a clinical trial of a specific drug) or the outcome (as would be the case when treating hypertension with the intention of gradually decreasing the blood pressure using a certain class of drugs). An early example of an intention-oriented language is the **Asbru** language [Shahar et al., 1996; Miksch et al., 1997; Seyfang et al., 2000], which is used within the **Asgaard** project [Shahar et al., 1998].

The concept underlying intention-based design is that plans, and therefore also therapy plans, are inextricably associated with intentions, and vice versa, as has been shown in detail in Bratman's comprehensive book on the subject [Bratman, 1987]. Access to the original process and outcome intentions of the guideline designers supports forming an automated critique of *where*, *when*, *how much* the care provider seems to be deviating from the suggested process of applying the guideline, and in *what way* and *to what extent* the care provider's outcome intentions might still similar to those of the author's (e.g., she might be using a different process to achieve the same outcome intention). Using a principle of rationality we can start from the assumption that the care provider is not trying to harm the patient and is probably trying to treat her main problems, although not necessarily in the fashion dictated by the guideline. Thus, one view of effective quality assessment is that it mainly consists of searching for a reasonable *explanation* that tries to understand the care provider's rational by comparing it to the design rational of the guideline's author. That design rational needs to be explicitly captured in a set of process and outcome intentions, which can optionally exist at every level (e.g., component) of the guideline. Thus, for example, both false positive and false negative alarms (i.e., undeserved harsh critique or undetected misguided care) might be prevented during quality assessment time. That time, of course, might be either retrospective or at the point of care, when the care provider proposes her plan, just before issuing an order.

Thus, intelligent quality assessment of guideline application requires (1) awareness of the guideline author's intentions, (2) knowledge of the effects of different interventions (e.g., to recognize substitution of a anti-hypertension drug by a drug that has a similar effect), and (3) a set of general and guideline-specific revision strategies (e.g., the critiquing module must realize that, in general, stopping administration of a drug that has a negative effect on some clinical parameter, such as Hemoglobin level, is equivalent to administering a drug, or performing an action, that has a positive effect on that parameter, such as performing a blood transfusion).

Note that *intentions* are much more specific than general *themes* such as reducing mortality and morbidity; these themes cannot be monitored effectively during the lifetime of the guideline's application. Intentions also are not as specific as effects on low-level physiological mechanisms that might lead to reduction in morbidity and mortality we rarely have precise mathematical models for such a complex chain of events. Rather, intentions exist at an intermediate level that captures the guideline designer's constraints on the (temporal) patterns that should emerge from correctly following the guideline's actions and from achieving its expected (short-term) outcomes.

5. Summary

There is a clear need for effective guideline-support tools at the point of care and at the point of quality assessment, which will relieve the current information overload on both care providers and administrators, as well as improve the safety of clinical care. To be effective, these tools need to be grounded in the patient's record, must use standard medical vocabularies, should have clear semantics, must facilitate knowledge maintenance and sharing, and need to be sufficiently expressive to explicitly capture the design rational (process and outcome intentions) of the guideline's author, while leaving flexibility at application time to the attending physician and their local favorite methods.

REFERENCES

Barnes, M. and Barnett, G.O. (1995). An architecture for a distributed guideline server. In: R. M. Gardner, ed., *Proceedings of the Annual Symposium on Computer Applications in Medical Care (SCAMC-95) (New Orleans, LA*), 233–237, Hanley & Belfus.

Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press: Cambridge, MA. Republished in 1999 by CSLI Publications: Stanford, CA.

Fox, J., Johns, N., and Rahmanzadeh, A. (1998). Disseminating medical Knowledge: the PRO*forma* approach. *Artificial Intelligence in Medicine*, 14: 157-181.

Grimshaw, J.M. and Russel, I.T. (1993). Effect of clinical guidelines on medical practice: A systematic review of rigorous evaluations. *Lancet*, 342: 1317–1322.

Herbert, S.I., Gordon, C.J., Jackson-Smale, A., and Renaud Salis, J-L. (1995). Protocols for clinical care. *Computer Methods and Programs in Biomedicine* 48: 21–26.

Hripcsak, G., Ludemann, P., Pryor, T.A., Wigertz, O.B., and Clayton, P.D (1994). Rationale for the Arden Syntax. *Computers and Biomedical Research*, 27: 291–324.

Johnson PD, Tu SW, Booth N, Sugden B, and Purves IN (2000). Using scenarios in chronic disease management guidelines for primary care. In Overhage M.J.,, Ed., *Proceedings of the 2000 AMIA Annual Symposium (Los Angeles, CA, 2000*), Hanley & Belfus, Philadelphia.

Karras BT, Nath SD, and Shiffman RN 2000). A Preliminary Evaluation of Guideline Content Mark-up Using GEM—An XML Guideline Elements Model. In Overhage M.J.,, Ed., *Proceedings of the 2000 AMIA Annual Symposium (Los Angeles, CA, 2000*), 942, Hanley & Belfus, Philadelphia.

Miksch, S., Shahar, Y., and Johnson, P. (1997). Asbru: A task-specific, intention-based, and time-oriented language for representing skeletal plans. In: *Proceedings of the Seventh Workshop on Knowledge Engineering Methods and Languages* (KEML-97) (Milton Keynes, UK, 1997), 9-1–9-20.

Miksch, S., Kosara, R., Shahar, Y., and Johnson, P. (1998). AsbruView: Visualization of time-oriented, skeletal plans. In: T*he Fourth International Conference on Artificial Intelligence*

Planning Systems 1998 (AIPS-98) (Carnegie-Mellon University, Pittsburgh, Pennsylvania), AAAI Press, Menlo Park, CA, 11-18.

Miller, P.L. (1986). *Expert Critiquing Systems: Practice-Based Medical Consultation by Computer*, Springer-Verlag, New York, NY.

Musen M. A., Carlson R. W., Fagan L. M., and Deresinski S. C. (1992). T-HELPER: Automated Support for Community-Based Clinical Research. *Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care*, Washington, D.C., 719-723.

Musen, M.A., Tu, S.W., Das, A.K., and Shahar, Y. (1996). EON: A component-based approach to automation of protocol-directed therapy. *Journal of the American Medical Information Association* 3(6): 367–388. URL: http://smi-web.stanford.edu/pubs/SMI_Abstracts/SMI-96-0606.html .

Ohno-Machado L, Gennari JH, Murphy SN, et al. (1998). The guideline interchange format: a model for representing guidelines. *Journal of the American Medical Informatics Association* **5** 357-72.

Peleg M, Boxwala A. A., Omolola O., Zeng Q., Tu, S.W, Lacson R., Bernstam, E., Ash, N., Mork, P., Ohno-Machado, L., Shortliffe, E.H., and Greenes, R.A. (2000). GLIF3: The Evolution of a Guideline Representation Format In Overhage M.J.,, Ed., *Proceedings of the 2000 AMIA Annual Symposium (Los Angeles, CA, 2000*), Hanley & Belfus, Philadelphia.

Seyfang S., Kosara R., and Miksch S. (2000). Asbru's Reference Manual, Asbru Version 7.2, Document Revision 0. *Vienna University of Technology, Institute of Software Technology, Vienna, Technical Report, Asgaard-TR-2000-3*. The Vienna Asgaard Project URL: http://www.ifs.tuwien.ac.at/~silvia/projects/asgaard/

Shahar, Y., Miksch, S., and Johnson, P. (1996). An intention-based language for representing clinical guidelines. In Cimino J.J., Ed., *Proceedings of the 1996 AMIA Annual Symposium (Washington, D.C., 1996*), 592-596, Hanley & Belfus, Philadelphia.

Shahar, Y., Miksch, S., and Johnson, P. (1998). The Asgaard project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artificial Intelligence in Medicine* (14): 29-51. URL: http://smi-web.stanford.edu/pubs/SMI_Abstracts/SMI-97-0702.html .

Sherman, E.H., Hripcsak, G., Starren, J., Jender, R.A., and Clayton, P. (1995). Using intermediate states to improve the ability of the Arden Syntax to implement care plans and reuse knowledge. In: R. M. Gardner, Ed., *Proceedings of the Annual Symposium on Computer Applications in Medical Care* (SCAMC-95) (New Orleans, LA, 1995), 238–242, Hanley & Belfus.

Shiffman, R.N., and Greenes, R.A. (1994). Improving clinical guidelines with logic and decision-table techniques. *Medical Decision Making*, 245-254.

Shiffman RN, Karras BT, Agrawal A, Chen R, Marenco L, Nath S. (in press). GEM: a proposal for a more comprehensive guideline document model using XML. *Journal of the American Medical Informatics Association*.

Tang PC and Young CY (2000). ActiveGuidelines: Integrating Web-Based Guidelines with Computer-Based Patient Records. In Overhage M.J.,, Ed., *Proceedings of the 2000 AMIA Annual Symposium (Los Angeles, CA, 2000*), Hanley & Belfus, Philadelphia.

Tu, S.W., Kahn, M.G., Musen, M.A., Ferguson, J.C., Shortliffe, E.H., and Fagan, L.M. (1989). Episodic Skeletal-plan refinement on temporal data. *Communications of ACM* 32: 1439–1455.