Y. Shahar

Medical Informatics Research Center
 Department of Information Systems
 Engineering
 Ben Gurion University
 Beer Sheva, Israel

Synopsis

Knowledge-Based Systems: Enhancing the Quality of Care

Since the early days of medical informatics (when that term, in fact, did not exist as such), there has been a continuing interest in automating medical diagnosis. The clinical-diagnosis task is now considered by the knowledge-modeling community as a very complex, domain-specific version of the generic classification task, not unlike detection of anomalies in digital circuits. The diagnosis task had seemed to often be the first to be tackled enthusiastically by newcomers to the medical-informatics field, especially those from the more mathematical and computational sciences. However, the clinicians themselves did not always equally share that enthusiasm, feeling that *management* of patients, rather than simply *classifying* their initial problem, was the real issue. That feeling was indeed supported by formal and informal surveys regarding the information needs of physicians.

Over the past two decades, it has become increasingly clear that supporting clinical *therapy* and continuous *management*, and in particular, enhancing the *quality* of that therapy by multiple runtime *quality-assurance* and retrospective *quality-assessment* methods, is the major new frontier. The encounter of the overwhelming

majority of patients with their clinicians is not the first one. Thus, often the issue at stake is not to classify the patient as a diabetes type II patient, but rather to make the difficult decision, based on past clinical course and present clinical data, how to *manage* that patient. Most of the health-care costs are now spent on management of patients who suffer from chronic conditions such as cardiovascular diseases, diabetes, pulmonary diseases, and chronic infectious diseases (e.g., AIDS).

The four papers in the Knowledge Processing and Decision Support category present four different aspects of tackling the various facets of the clinical-management task:

- a. Increasing the use of preventive care in hospitalized patients, by using computerized reminders integrated within an order-entry system [Dexter et al., 2001];
- b. Analyzing in depth the relationship between the approach of sharing procedural clinical knowledge regarding continuous, long-term medical care, represented as clinical guidelines in the GLIF3 language, with that of using one-time reminders, represented as individual rules in the Arden syntax [Peleg et al., 2001];
- c. Increasing our insight and knowledge regarding the management of patients who have head injuries, by exploiting not only initial, “demographic,” data, but also accumulating, time-oriented clinical data; and by discussing the deeper meaning of these data with medical experts, using the structure of decision-trees induced automatically from a database of patients who have had head injuries [McQuatt et al., 2001]; and
- d. Continuously assessing the quality of surgical care, using a risk-adjusted cumulative sum method that quickly and graphically zeroes in on changes in surgical outcomes, thus potentially supporting remedial measures [Steiner et al., 2001].

Much of the major progress over the past several years in the task of supporting patient management has occurred in the area of automated support to guideline-based care. Thus, a brief overview of the state of the art in that area would be useful.

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 sche-

matic plans, at varying levels of abstraction and detail, for management over extended periods of patients who have a particular clinical condition (e.g., insulin-dependent diabetes). **Clinical protocols** are typically highly detailed guidelines, often used in areas such as oncology and experimental clinical trials. **Reminders** and **alerts** can be viewed as “mini guidelines”, useful mostly for representing a single rule that needs to be applied whenever the patient’s record is accessed, as opposed to representation of a long-term plan [Peleg et al., 2001]. Their effectiveness (as part of an automated system) in outpatient care has been demonstrated repeatedly, but the paper featured in this section demonstrates forcefully that they are highly effective (especially for promoting preventive care, such as pneumococcal vaccination) also in hospital environments [Dexter et al., 2001]. 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.

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 possible view, however, 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), that is, process (care-provider *action*) and outcome (patient *state*) **intentions** [Shahar et al., 1998]. 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.

Most clinical guidelines exist only in free-text format and are 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 than 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*.

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. Several simplified approaches to the task of supporting guideline-based care that *do* use the patient’s data encode guidelines as elementary state-transition tables or as situation-action rules dependent on the electronic medical record, as was attempted using the Arden syntax [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 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). However,

rule-based approaches, such as MLMs, typically do not include an intuitive representation of the guideline's clinical logic, have no semantics for the different types of clinical knowledge represented, lack the ability to easily represent and reuse guidelines and guideline components as well as higher, meta-level problem-solving knowledge, cannot represent intended ambiguity (e.g., when there are several options and several pro and con considerations, but no single action is, or should be, clearly prescribed) [Peleg et al., 2001], and do not support application of guidelines over extended periods of time, [Peleg et al., 2001] as is necessary to support the care of chronic patients. On the other hand, as Peleg et al. also point out, such approaches do have the advantage of simplicity when only a single alert or reminder is called for, and the heavier machinery of higher-level languages is uncalled for and might even be disruptive. Thus, they might be viewed as complementary to complex guideline representations.

During the past 20 years, there have been several efforts to support *complex* guideline-based care over time in automated fashion. 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], PROforma [Fox et al., 1998], the guideline interchange format (GLIF) [Ohno-Machado et al., 1998; Peleg et al., 2001], the European PRESTIGE project [Gordon and Veloso, 1996], and the British Prodigy project [Johnson et al., 2000].

Most of the approaches 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 [Miller, 1986], have used a *critiquing* approach, in which

the physician suggests a specific therapy plan and gets feedback from the program. The Asgaard project [Shahar et al., 1998] uses the *Asbru* language, which supports both an expressive, time-oriented, prescriptive specification of recommended interventions, and a set of meta-level annotations, such as process and outcome intentions of the guidelines, which support also a critiquing approach for retrospective quality assessment. Access to the original process and outcome intentions of the guideline designers supports forming an automated critique of *where*, *when*, and by *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 be similar to those of the author's (e.g., she might be using a different process to achieve the same outcome intention). Thus, effective quality assessment includes 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. (It is perhaps a specific instance of a rather general observation, that critiquing an agent's actions must always include at least an attempt to understand that agent's reasons for such actions).

Other recent approaches to support guideline use at the point of care enable a Web-based connection from an electronic patient record to an HTML-based set of rules, such as is done in the ActiveGuidelines model [Tang and Young, 2000], which is embedded in a commercial electronic medical record system. However, such approaches have no standardized, sharable, machine-readable representation of guidelines that can support multiple tasks such as automated application and quality assurance, and are not intended for representation of complex care plans over time. A recent

framework, GEM, enables structuring of a text document containing a clinical guideline as an *extensible markup language (XML)* document, using a well-defined XML schema [Shiffman et al., 2000]. However, GEM is an application running on a stand-alone computer, and the framework 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. Thus, it seems that the future lies with architectures that support the full life cycle, from guideline specification by experts, through a computable representation, to a locally customized guideline; GLIF3 is one of the architectures supporting such a life cycle [Peleg et al., 2001].

In summary, there is a clear need for effective guideline-support tools at the point of care and at the point of critiquing, which will relieve the current information overload on both care providers and administrators. 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 physicians and their local favorite methods.

References

- Barnes M, Barnett GO. An architecture for a distributed guideline server. In: Gardner RM, editor. Proceedings of the Annual Symposium on Computer Applications in Medical Care (SCAMC-95) (New Orleans, LA). Philadelphia: Hanley & Belfus; 1995. p. 233-7.
- Dexter PR, Perkins S, Overhage MJ, Maharry K, Kohler RB, McDonald CJ. A computerized reminder system to increase the use of preventive care for hospitalized patients. *N Engl J Med* 2001; 345(13): 965-70.

- Fox J, Johns N, Rahmzadeh A. Disseminating medical Knowledge: the PROforma approach. *Artif Intell Med* 1998; 14:157-81.
- Gordon C, Veloso M. The PRESTIGE Project: Implementing Guidelines in Healthcare. *Medical Informatics Europe '96*: IOS Press. 1996. p. 887-91.
- Grimshaw JM, Russel IT. Effect of clinical guidelines on medical practice: A systematic review of rigorous evaluations. *Lancet* 1993; 342:1317-22.
- Herbert SI, Gordon CJ, Jackson-Smale A, Renaud Salis J-L. Protocols for clinical care. *Comput Methods Programs Biomed* 1995; 48:21-2.
- Hripcsak G, Ludemann P, Pryor TA, Wigertz OB, Clayton PD. Rationale for the Arden Syntax. *Comput Biomed Res* 1994; 27: 291-324.
- Johnson PD, Tu SW, Booth N, Sugden B, Purves IN. Using scenarios in chronic disease management guidelines for primary care. In: Overhage MJ, editor. *Proceedings of the 2000 AMIA Annual Symposium* (Los Angeles, CA, 2000). Philadelphia: Hanley & Belfus; 2000.
- Miller PL. *Expert Critiquing Systems: Practice-Based Medical Consultation by Computer*. New-York, NY: Springer-Verlag; 1986.
- McQuatt A, Sleeman D, Andrews PJD, Corruble V, Jones PA. Discussing anomalous situations using decision trees: A head injury case study. *Methods Inf Med* 2001; 40 (5):373-9.
- Musen MA, Carlson RW, Fagan LM, Deresinski SC. T-HELPER: Automated Support for Community-Based Clinical Research. *Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care*, Washington, D.C.; 1992. p. 719-23.
- Musen MA, Tu SW, Das AK, Shahar Y. EON: A component-based approach to automation of protocol-directed therapy. *J Am Med Inform Assoc* 1996; 3(6): 367-88.
- Ohno-Machado L, Gennari JH, Murphy SN, Jain NL, Tu SW, Oliver DE, et al. The guideline interchange format: a model for representing guidelines. *J Am Med Inform Assoc* 1998; 5:357-72.
- Peleg M, Boxwala A, Bernstam E, Tu SW, Greenes R, Shortliffe EH. Sharable representation of clinical guidelines in GLIF: Relationship to the Arden syntax. *J Biomed Inform* 2001; 34:170-81.
- Shahar Y, Miksch S, Johnson P. The Asgaard project: A task-specific framework for the application and critiquing of time-oriented clinical guidelines. *Artif Intell Med* 1998; (14): 29-51.
- Sherman EH, Hripcsak G, Starren J, Jender RA, Clayton P. Using intermediate states to improve the ability of the Arden Syntax to implement care plans and reuse knowledge. In: Gardner RM, editor. *Proceedings of the Annual Symposium on Computer Applications in Medical Care* (SCAMC-95) (New Orleans, LA, 1995), Philadelphia: Hanley & Belfus; 1995. p. 238-42.
- Shiffman RN, Karras BT, Agrawal A, Chen R, Marengo L, Nath S. GEM: a proposal for a more comprehensive guideline document model using XML. *J Am Med Inform Assoc* 2000; 7(5):488-98.
- Steiner SH, Cook RJ, Farewell VT. Risk-adjusted monitoring of binary surgical outcomes. *Med Decis Making* 2001; 21:163-9.
- Tang PC, Young CY. ActiveGuidelines: Integrating Web-Based Guidelines with Computer-Based Patient Records. In: Overhage MJ, editor. *Proceedings of the 2000 AMIA Annual Symposium* (Los Angeles, CA, 2000), Philadelphia: Hanley & Belfus; 2000.
- Tu SW, Kahn MG, Musen MA, Ferguson JC, Shortliffe EH, Fagan LM. Episodic Skeletal-plan refinement on temporal data. *Commun ACM* 1989; 32:1439-55.

Address of the author:
 Yuval Shahar M.D., Ph.D.
 Medical Informatics Research Center
 Department of Information
 Systems Engineering
 Ben Gurion University
 Beer Sheva, Israel 84105
 E-mail: yshahar@bgumail.bgu.ac.il