Why Clinical Decision Support is Hard to Do

Robert A. Greenes, M.D., Ph.D.
Brigham and Women's Hospital, Harvard Medical School, Boston, MA

The goals of improving healthcare quality, safety, and cost-effectiveness have stimulated great interest in computer-based clinical decision support (CDS). A number of projects have been singularly successful in demonstrating effectiveness of CDS, and a growing number of products have built-in CDS capabilities. Despite this, the harsh reality is that successes have not been easily replicated, the appearance in products is spotty, and there is no consensus on how to deploy CDS broadly to achieve its promised goals.

The root problem is that CDS is both (a) deceptively easy for one-off implementations, and (b) extremely hard to do robustly. The apparent simplicity of an *if…then* rule, as in an alert or reminder, for example, is seductive. Potential users are blissfully unaware of the significant infrastructure needed to create libraries of such rules, maintain/update them, and deploy them. This is even more daunting when one considers a clinical enterprise with a multi-vendor information system environment, or the problem of sharing knowledge libraries across enterprises or under some neutral aegis such as a professional specialty society.

A key requirement in order to address those more ambitious targets for deployment and use is the need to create knowledge representations that are sufficiently transparent to be understood by domain experts, computationally unambiguous, and capable of being readily interfaced with, or adapted to working in, a variety of platforms. For a decision rule, this involves the ability for communication to occur between an application in the host system and the logical evaluation process or service, in order (a) to initiate the evaluation and pass necessary patient data; and also, (b) once an action recommendation has been determined, to inform the application. The first of these tasks involves mapping of references to data elements, in a generic information model, in a decision rule to/from the host-specific way it is referenced as a term, and to the data model corresponding to how its values are stored in the host database. The second task involves translating an action recommendation to the host services that will carry it out.

To deliver CDS in a coordinated way on an enterprise-wide basis, (or beyond that, to multiple sites), scales up this challenge significantly. One needs to recognize that there are three overlapping lifecycle processes involved. First, any item of knowledge goes through a lifecycle, from being generated, validated, converted to a formal computable representation, and stored, to being updated.

Second, optimal use of a knowledge item depends on it being delivered effectively, in terms of its interface with clinical applications. How to do this is typically learned empirically by initial implementation, pilot testing, and successive refinement – a lifecycle process. Issues that must be addressed include ease of integration into applications, time and effort required for use, quality of the advice, and perception of value. This lifecycle process relates not only to what is needed for initial success but being able to replicating the success, become a product, become profitable, and survive (and eventually to require updates) .

The third lifecycle process involves managing knowledge bases, disseminating them, and keeping track of where the knowledge is used. This lifecycle deals with the corpus of knowledge, not individual knowledge elements, in contrast to the first lifecycle described above. New or revised knowledge items must be reviewed in the context of other knowledge being used, to identify overlaps, inconsistencies, etc. The knowledge must be organized and indexed for retrieval, and the places where it is used tracked. When knowledge is updated, those places must be found so that the necessary changes can be determined and carried out.

The above lifecycles are all complex undertakings, each involving different kinds of participants. Knowledge generation/validation relies on evidence-based medicine experts and is based on clinical research, meta-analysis, and other methods. Computer system design/interface involves system developers, experts in human computer interaction, workflow modelers, and clinician. Knowledge management involves standards organizations, content management/collaboration experts, and other technical individuals, as well as those who can articulate the business case and can commit organizations to align their resources and activities to undertake the effort.

Widespread deployment of CDS is thus fraught with not only scientific and methodological issues, but with cultural, sociological, regulatory, and business challenges. To make progress, it is essential that sufficient attention and resources be devoted to these processes, to learn how to optimize them, leverage efforts of the community, build the necessary infrastructure, and figure out how to pay for it.

Robert A. Greenes, MD, PhD
Brief biosketch

Dr. Greenes directs the Decision Systems Group, a Harvard-based biomedical informatics research and development laboratory at the Brigham and Women's Hospital, The DSG pursues informatics approaches for data mining and knowledge discovery, knowledge management, decision support, natural language processing, and education, with applications in biomedical research, clinical systems, image-based systems, consumer health, and public health systems. With M.D. and Ph.D. in applied mathematics/computer science from Harvard, Dr. Greenes is Board Certified in Diagnostic Radiology. He is the Distinguished Chair in Biomedical Informatics at the Brigham and Women's Hospital, Professor of Radiology and Health Sciences and Technology at Harvard Medical School, and Professor of Health Policy and Management, Harvard School of Public Health. He is Program Director of the Boston Research Training Program in Biomedical Informatics, based at the Harvard-MIT Division of Health Sciences and Technology. Dr. Greenes is author/editor of a forthcoming book entitled "Clinical Decision Support: Computer-Based Approaches to Improving Healthcare Quality and Safety", to be published by Elsevier. He is a Member of the Institute of Medicine, and Fellow of the American College of Medical Informatics (also its past President), the American College of Radiology, and the Society of Computer Applications in Radiology.

The primary reference for the presentation is: Greenes RA (ed.). *Clinical Decision Support: Computer-Based Approaches to Improving Healthcare Quality and Safety*, to be published by Elsevier, 2006 or early 2007. The book contains a number of contributed chapters by ACMI fellows.