

EXPERIENCE REPORT

Knowledge management in the era of digital medicine: A programmatic approach to optimize patient care in an academic medical center

Jane L. Shellum¹ | Rick A. Nishimura² | Dawn S. Milliner² | Charles M. Harper Jr.² | John H. Noseworthy²

¹Information Technology, Knowledge and Delivery Center, Mayo Clinic, Rochester, Minnesota

²Mayo Clinic College of Medicine, Mayo Clinic, Rochester, Minnesota

Correspondence

Jane L. Shellum, Section Head in Information Technology and Administrator of the Knowledge and Delivery Center, Mayo Clinic, 200 1st St SW, Rochester, MN 55905. Email: shellum.jane@mayo.edu

Abstract

Introduction The pace of medical discovery is accelerating to the point where caregivers can no longer keep up with the latest diagnosis or treatment recommendations. At the same time, sophisticated and complex electronic medical records and clinical systems are generating increasing volumes of patient data, making it difficult to find the important information required for patient care. To address these challenges, Mayo Clinic established a knowledge management program to curate, store, and disseminate clinical knowledge.

Methods The authors describe AskMayoExpert, a point-of-care knowledge delivery system, and discuss the process by which the clinical knowledge is captured, vetted by clinicians, annotated, and stored in a knowledge content management system. The content generated for AskMayoExpert is considered to be core clinical content and serves as the basis for knowledge diffusion to clinicians through order sets and clinical decision support rules, as well as to patients and consumers through patient education materials and internet content. The authors evaluate alternative approaches for better integration of knowledge into the clinical workflow through development of computer-interpretable care process models.

Results Each of the modeling approaches evaluated has shown promise. However, because each of them addresses the problem from a different perspective, there have been challenges in coming to a common model. Given the current state of guideline modeling and the need for a near-term solution, Mayo Clinic will likely focus on breaking down care process models into components and on standardization of those components, deferring, for now, the orchestration.

Conclusion A point-of-care knowledge resource developed to support an individualized approach to patient care has grown into a formal knowledge management program. Translation of the textual knowledge into machine executable knowledge will allow integration of the knowledge with specific patient data and truly serve as a colleague and mentor for the physicians taking care of the patient.

KEYWORDS

computer-interpretable guidelines, knowledge management, knowledge representation

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2017 The Authors. Learning Health Systems published by Wiley Periodicals, Inc. on behalf of the University of Michigan

1 | INTRODUCTION

The creation and dissemination of medical knowledge are of critical importance in today's health care systems. The medical world is in the midst of a knowledge explosion driven by constant advances in diagnostics and treatments as well as the intersection of care delivery with genomics, proteomics, and metabolomics. While this whirlwind of information stands to further improve a patient's health and well-being, the pace of discovery has accelerated to a point where it is no longer possible for caregivers to keep up. It has been estimated that each day, over 1500 new journal articles and 55 new clinical trials are indexed in the National Library of Medicine Medline database.¹ Less than 1% of published clinical information is likely to be relevant for a particular physician; yet that 1% may offer lifesaving information for an individual patient.² All these factors now contribute to the knowledge overload, which all practicing physicians face in providing optimal care for their patients.

Mayo Clinic provides multispecialty, interdisciplinary care of patients with complex medical and surgical problems using an integrated team that focuses on all aspects of patient care. From the early 1900s, when Henry Plummer introduced the shared medical record, Mayo Clinic has emphasized shared clinical knowledge as a force integrating multiple disciplines around the care of an individual patient. As it entered the era of digital medicine, Mayo Clinic recognized that new solutions would be required to (1) perpetuate its history of generating new knowledge, (2) vet and integrate that which is learned by others, and (3) actively manage this clinical knowledge to bring it immediately and seamlessly into the clinical practice. Thus, the knowledge management program was established with the responsibility to curate, store, and update Mayo-vetted clinical knowledge into a single repository. The following outlines the development of the knowledge

management program, its role in the Mayo practice, its efforts to integrate clinical knowledge into the workflow, and the future vision for the program.

2 | BACKGROUND

The clinical knowledge applied to patient care is based on the synthesis of clinical experience, in-depth understanding of diagnostic testing and therapies, and critical analysis of clinical trials examining the effect of a drug or intervention. There are multiple knowledge sources, ranging from textbooks to medical journals to online medical resources, but controversies and differing opinions always exist among physicians. Mayo Clinic has specialty and subspecialty experts who share their knowledge with colleagues either through formal consultation or, just as often, through informal conversations in which colleagues provide quick answers to focused questions about patient care. These encounters are viewed as a "source of truth" for questions about patient care. However, the rapid growth in the number of physicians and scientists and continued subspecialization has made it more difficult for staff members to know who might have the expertise to answer their questions. In 2006, leadership summarized this growing challenge with the question, "does Mayo know what Mayo knows?" (Figure 1).

3 | ASKMAYOEXPERT—A POINT-OF-CARE RESOURCE

In response to this challenge, Mayo Clinic created an online point-of-care resource called AskMayoExpert (AME). The purpose of AME is to provide the clinician with Mayo-vetted clinical knowledge at

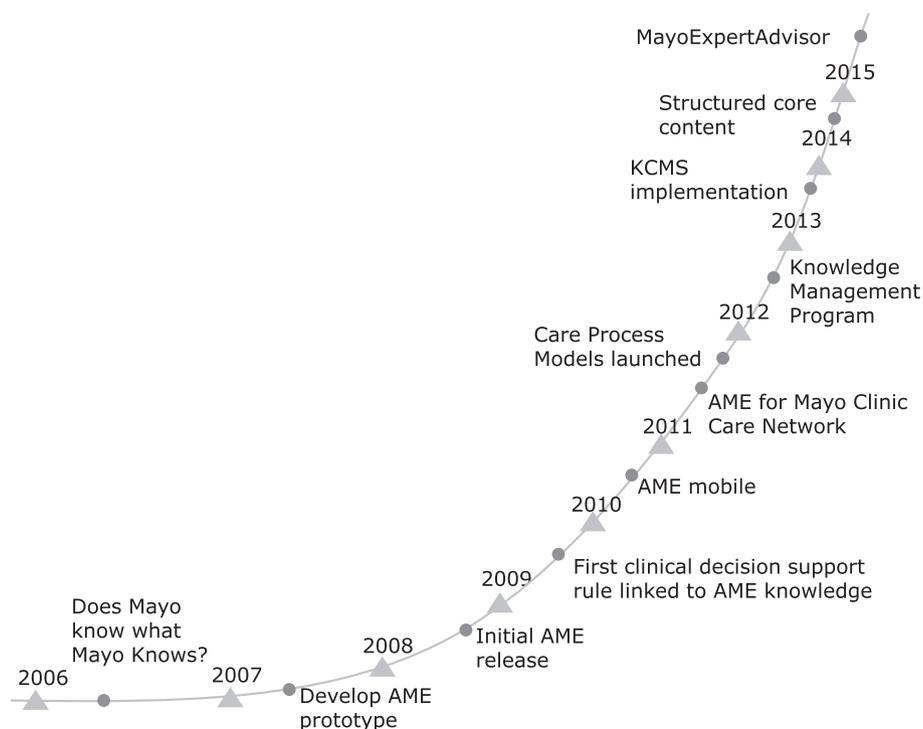


FIGURE 1 Knowledge Management time line. This figure illustrates the major milestones in the development of the Knowledge Management program at the Mayo Clinic

the point of care. AskMayoExpert was developed based on the concept of gist and verbatim memory and learning. Verbatim memories focus on the “surface forms” of information, that is, a series of facts, while gist memory is about the meaning and interpretation of the facts. A point-of-care tool is most effective for clinicians who understand the gist but require assistance with keeping up with all of the verbatim information that relates to the gist.^{3,4} AskMayoExpert provides concise, relevant, and clinically applicable answers to clinical questions, assuming an existing knowledge of the “gist” of medical decision making. For example, a clinician understands the “gist” that it is critical to stop anticoagulation before a procedure with a high-bleeding risk but a point-of-care tool can provide the concise, actionable answer in the safest duration of cessation of an anticoagulant drug prior to a procedure.

Experts were asked to compile their most frequently asked questions (FAQs) from colleagues and generate clinically relevant responses. These responses were stored in a database annotated with Systematized Nomenclature of Medicine terms to improve search accuracy. AskMayoExpert also developed a database in which physicians would declare their specific areas of expertise, again, using the Systematized Nomenclature of Medicine taxonomy. This created a mechanism for managing increasing complexity, so that if a patient care question is not answered by an FAQ, the physician can identify and contact an expert. If users are looking for more in-depth, encyclopedic information, AME can also pass search terms through to other commonly used resources such as UpToDate, Access Medicine, or Mayo Libraries.

An initial version of the application was released beginning in early 2009. Over the next 2 years, the application and content were iteratively enhanced based on feedback from users. In the fall of 2010, the application and content were deemed ready for broader release, and a communication campaign was launched to increase awareness of the application. Utilization has continued to grow (Figure 2), with over 80% of Mayo staff having used the application. Research has shown that AME is of high clinically relevant value to the users.⁴ Although initially targeted at generalist physicians, the application has been widely adopted by specialists, residents, mid-level providers, and nurses as well.

The greatest challenge in building the AME system was developing a process for creation and capture of clinical knowledge that would assure its credibility and acceptance. Subject matter experts, identified by practice leadership, work with medical writers and a standard

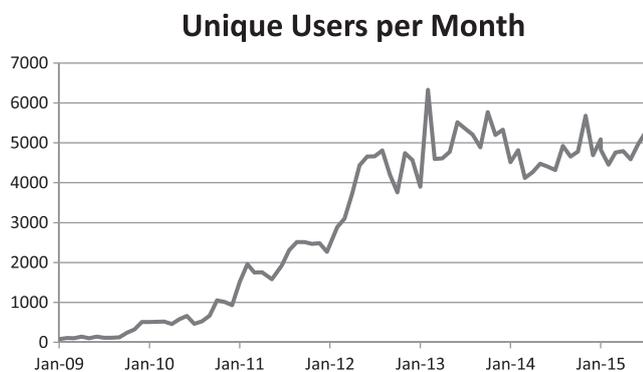


FIGURE 2 AskMayoExpert utilization growth. This figure illustrates the increase in unique users per month since the introduction of AME

interview process to develop the AME content. The content is then evaluated and vetted by knowledge content boards (KCBs), a select group of highly recognized clinicians and educators from each department or division. There are now 44 KCBs representing a variety of medical and surgical specialties and subspecialties. These boards are responsible not only for vetting the FAQs but also for responding to user feedback and rapidly incorporating new information regarding tests and treatments. Under their leadership, the content has grown steadily and now comprises over 12 000 individual pieces of content, or “knowledge bytes,” covering more than 1500 topics (Figure 3). All content is reviewed every 6 to 12 months to assure that it remains current. This level of review requires a significant time commitment from physicians. The institutional leadership has provided the members of the KCBs with dedicated time to review and update the knowledge on an ongoing basis, indicative of the value the institution places on the knowledge management. Participation on the KCBs is recognized as an academic contribution by the Mayo Academic Appointments and Promotions Committee.

4 | CORE CLINICAL CONTENT—A FOUNDATION OF KNOWLEDGE MANAGEMENT

This content created for AME is now considered as “core clinical content” and has become the center of our knowledge management program. To better manage this content, we invested in a centralized knowledge management system, referred to as the knowledge content management system (KCMS), using Sitecore for the management of knowledge content and TopBraid for the management of ontologies. The clinical content generated for AME is divided into sections using Sitecore templates, which include specific concepts such as diagnosis, treatment, prevention, and follow-up. Each section is manually annotated by a trained ontologist, with annotation properties for subject, secondary subject focus, audience, and person group. These annotations provide rich descriptive metadata, and plans are underway

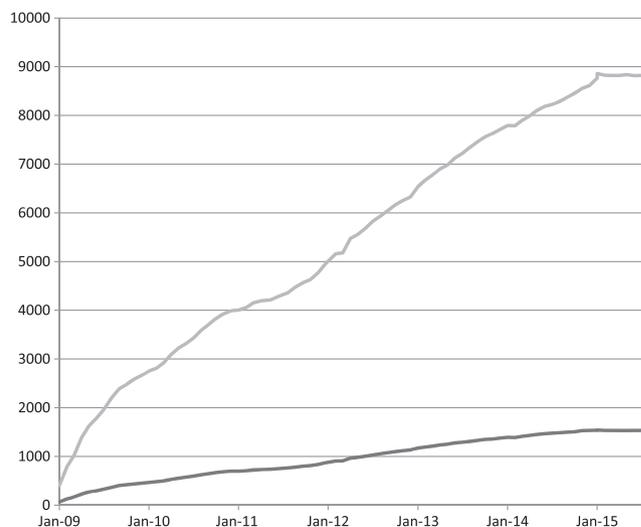


FIGURE 3 AskMayoExpert content growth. This chart shows the increase over time in the numbers of topics and frequently asked questions housed in AME

to enhance the KCMS to more fully leverage the annotations both for delivery and for the management of the knowledge. These sections are stored in an XML format and dynamically delivered through web pages, applications built on Sitecore including AME, or application program interfaces to other systems.

The core clinical content serves as the basis for text-based derivatives such as patient education materials and consumer health information. In addition, protocols, order sets, alerts, and reminders are developed based on the core clinical knowledge. These knowledge artifacts are cataloged in the KCMS and linked to the core clinical knowledge from which they are derived. This streamlines the process for capturing and vetting expert knowledge and ensures that all the clinical content is consistent and reflects Mayo's combined clinical knowledge. Any change or update in the core clinical knowledge is rapidly incorporated into all audience-specific channels for dissemination (Figure 4).

5 | CARE PROCESS MODELS— STANDARDIZATION OF BEST PRACTICES

Mayo Clinic emphasizes standardization of best practices. The practice is organized into specialty councils, consisting of clinical leaders in all specialties throughout the enterprise. These specialty councils are charged with identifying best practices based on both evidence and the consensus of experts, to be used as a basis for diagnosis and treatment of medical conditions; the AME team was charged with developing a mechanism to represent them and make them easily findable, understandable, and actionable at the point of care.

The care process model's (CPM) format was designed to guide a clinician through the care of a patient with a particular disease or disorder, providing concise, actionable care recommendations for both optimal patient management and point-of-care education. The CPMs are organized into a flow of decision steps and action steps. Each step in the CPM algorithm expands to provide more detailed practical information such as specific dosing and titration schedules, ordering

instructions, patient education materials, and teaching points. This additional information may include not only text but also external links, interactive calculators, or video. The clear, concise, and actionable language used in the CPMs is intended to encourage their adoption and application.⁵

6 | INTEGRATION OF CLINICAL KNOWLEDGE INTO THE WORKFLOW

The initial functionality of AME required users to launch the application and search for answers to their clinical questions. Navigation was simplified by embedding links to the application on the Mayo intranet home page, practice websites, and within the electronic medical record (EMR). With the introduction of the meaningful use requirements of the HITECH Act,⁶ electronic health records (EHRs) began to offer "infobutton" functionality to provide access to relevant knowledge resource, based on the clinical context provided by data in the EHR.⁷ Mayo's EMR's infobutton is configured to retrieve content from AME. These efforts have streamlined navigation to AME, but to fully apply, clinical knowledge requires that the knowledge be individualized and integrated into the clinical workflow. The MayoExpertAdvisor (MEA) application is being developed to meet this need. The CPMs are converted into executable rules, which leverage patient data, both structured and unstructured, to present patient-specific care recommendations within Mayo's home-grown EMR viewer. The care recommendations are presented along with the supporting data and any relevant calculations and risk scores. Risk scoring tools are prepopulated with patient data, and providers can alter the displayed data to do "what if" scenarios without changing the underlying values in the EMR. The implementation approach is nontransactional; that is, rather than having event-triggered recommendations or actions presented to the clinician, the CPMs are evaluated for any applicable recommendations at the time the chart is opened, and these recommendations are available to the care giver when needed during the encounter. A visual indicator in the navigation bar shows that there

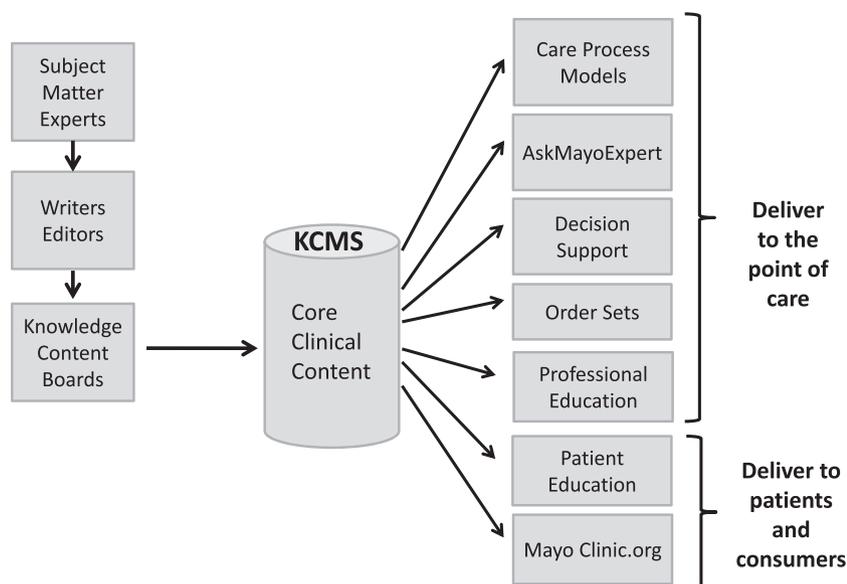


FIGURE 4 Knowledge Management at Mayo Clinic. This diagram illustrates the process by which subject matter experts, working with writers and editors, generate core clinical content, which is vetted by Knowledge Content Boards and stored in the Knowledge Content Management System. This content serves as the basis for a variety of mechanisms for delivering knowledge to providers, patients, and consumers

is a recommendation for the patient and the clinician can navigate to the MEA page to see it at any time. MayoExpertAdvisor is currently being evaluated in a randomized controlled trial in the primary care practice at one site.

The current process for converting the CPMs into the recommendations in MEA is as follows:

6.1 | Knowledge representation

While the CPMs represent an algorithmic approach to management of a condition, they are not sufficiently structured to enable the direct extraction of an executable algorithm. Therefore, a knowledge engineer “deconstructs” the CPM into an if/then format, similar to Arden Syntax, to create an unambiguous representation of the logic to be used by the software developer to write the executable rules.

One of the advantages to this approach is that Arden Syntax, first published as an HL7 standard in 1999, is one of the earliest and most widely used standards for representing clinical knowledge in an executable format and is relatively easily understood by subject matter experts. With respect to modeling guidelines, however, the use of Arden Syntax has limitations. Arden is fundamentally made up of independent medical logic modules that do not support the task network model (telecommunications management network) in which a network of tasks unfolds over time.⁸ In addition, the medical logic module approach to Arden Syntax is centered on individual event-condition-action (ECA)-type rules, best suited for alerts and reminders. It does not easily support process flow or grouping of decisions nor does it easily support nondeterministic decisions.

6.2 | Data specification

For each proposition or input to the rules, the specific data elements must be defined. This is done through identification of value sets using standard terminologies (RxNorm, LOINC, and ICD-10) and defining natural language processing algorithms. These must in turn be mapped to each of the 3 EMRs in use at Mayo. The value sets are physician vetted and managed by Mayo's terminology team. Data specification also addresses process measurement; as each CPM is analyzed, the specific process metrics and the data elements needed for each are defined.

While this process ensures that the rules running in MEA are a faithful reproduction of the original, it has shortcomings. The process is complex and labor intensive, and the execution of the full CPM is incomplete. Any given executable CPM is made up of a number of interrelated knowledge assets such as rules, calculations, scales and scoring models, value sets, and natural language processing algorithms, each of which is potentially reused by other CPMs and other knowledge delivery applications and which must therefore be managed individually. In addition, except for the use of standard terminologies for the data definitions, the current approach is not standards based and does not allow for potential sharing of the executable versions.

In seeking a more robust, scalable approach, we reviewed the literature on computer-interpretable guidelines (CIGs). Although the CPM format was developed internally to meet specific organizational

goals, CPMs and guidelines are similar in structure and intent. The Institute of Medicine defines guidelines as “systematically developed statements to assist practitioner and patient decisions about appropriate health care for specific clinical circumstances.”⁹ Like guidelines, CPMs are systematically developed and focused on clinical decisions for specific conditions. More important in considering the applicability of CIG standards to the CPM process, they share many of the same structural characteristics as CIGs. A review of CIG models describes components that are shared across models¹⁰ Care process models are built using a home-grown authoring tool, and their components map to existing models as follows:

CPM Component	Definition	CIG Primitive
Step	Content describing actions to be taken (eg, order tests or examine patient)	Action
Decision	Branch point based on patient criteria (eg, findings or risk scores) with 2 possible alternatives	Decision
Decision choice	Describes the possible paths to a subsequent step or decision (generally yes/no)	Decision
Branch	Describes multiple paths, any one of which can be taken	Decision
Branch choice	Describes the criteria for each path (eg, risk score > 3)	
Link to external CPM	Provides navigation to a CPM, which could be considered a subset of the current CPM (eg, hypoglycemia management within diabetes CPM)	Nested guideline

Besides the components, CPMs share other characteristics with CIGs. First, as the name implies, the CPMs represent the process of care. They have scheduling constraints, that is, they include a sequence in which decisions and actions should occur. Second, they include the notion of nested guidelines. For example, the CPM for inpatient management of diabetes includes links to CPMs for management of hypoglycemia and ketoacidosis. Third, the CPMs include the concept of a patient state—for example, the patient requires an urgent cardioversion and has a therapeutic international normalized ratio that is the patient state within the CPM that informs the decision of whether a transesophageal echocardiogram is required. Finally, the CPMs include the patient data needed to make any given decision. Although the data elements are listed only as text, they provide a starting point toward understanding the clinical concepts needed to execute the CPM.

Attempts to develop CIGs began in the 1990s. The efforts were driven by the potential of guidelines to improve health care by modeling medical knowledge, driving clinical decision support efforts, monitoring the care processes, supporting clinical workflows and anticipating resource requirements, serving as a basis for training through simulation, and conducting clinical trials.¹¹ However, it is precisely this broad range of possible benefits that made it challenging to create one model that would serve every situation.¹² As a result, there have been many attempts at formalization of guidelines, and while some are in limited clinical use, many are still largely academic undertakings.

Mechanisms to share or reuse CIGs seek to maximize the benefit and facilitate the broad implementation of guidelines.¹³ GLIF3 is an example of formal guideline representation that was developed to enable the sharing of guidelines across institutions. GLIF3 is designed with the flexibility necessary to express guidelines for a variety of scenarios, including screening, diagnosis, and treatment, for acute and chronic problems, in primary and specialty care. While the GLIF model itself does not yield a fully executable guideline, work has been done to combine it with GELLO as an expression language and GLEE as an execution engine.¹⁴

Another approach to re-using CIGs is the service oriented approach. An example of this approach is SEBASTIAN, which uses web services to submit patient data and return clinical decision support results. The goal of this work was to provide “write once, run anywhere” functionality, while supporting ease of authoring in an understandable framework.¹⁵

The SAGE project, in which Mayo Clinic participated, was specifically focused on integrating guidelines into commercial clinical systems. Although it adopted many of the features of other models (activity graphs from EON and GLIF3, decision maps from PRODIGY, and decision model from PROforma),¹⁶ SAGE specifically focused on context, including triggering events, roles, resources, and care settings.¹⁷ The approach examined EHR-specific workflows and looked for opportunities for clinical decision support interventions. In particular, SAGE invoked context-appropriate order sets as a clinical decision support intervention. This ambitious project introduced new concepts into the guideline modeling discussion, which exposed advantages and disadvantages. The close integration with workflow has the potential to optimize the user experience by presenting the right intervention to the right person at the right time but, at the same time, requires more maintenance and updating of guidelines for changes in workflows and limits interoperability.¹⁶

Quaglini et al describe another workflow-focused approach to implementing clinical guidelines, which combines a formal representation of the medical knowledge with an organizational ontology, which describes agents, roles, resources, and tasks to model and implement “care flows.” This work provides an example of “separation of concerns” in which the medical knowledge and workflow knowledge are maintained separately to improve flexibility and ease of maintenance.¹⁸

Each of these modeling approaches has shown promise. However, because each of them addresses the problem from a different perspective, there have been challenges in coming to a common model for CIGs. Because of this, more recent approaches have focused on breaking down into components and focusing on the standardization of these components, deferring the orchestration.¹² Given the current state of guideline modeling and the need for a near-term solution, this is the approach that Mayo Clinic will likely take for executable CPMs.

7 | KNOWLEDGE MANAGEMENT: FUTURE DIRECTION

The future direction of the knowledge management program will focus both on continued exploration of models for representing CPMs and increasing our focus on measuring the impact of our work.

Additional exploration of model-driven knowledge-based tools to support clinical reasoning and decision making is in its early stages. The CPM could be represented as a decision-action model, where for each decision, a set of inputs define the patient data needed for the clinician to make the decision and a set of actions (generally orders) are offered as outputs. The decision itself is left to the clinician, but the summarization and presentation of the relevant data, along with brief narrative guidance, reduce the cognitive load. This approach is grounded in human-computer interaction principles, which stress the importance of external representation in distributed cognition.¹⁹ The approach is further informed by informatics research that has addressed the challenge of fully describing the context of a patient situation. This model has been referred to as a “GPS” model because it provides clinicians with relevant information about their current position and, given a destination or goal, can provide guidance to reach the destination. Providing full context for a decision maker requires an understanding of the disease process, the care process, the workflow process, and the information that describes each. An important facet of this approach is the role of situation awareness in the clinical decision-making process. Situation awareness combines an individual's perception and comprehension of a dynamic environment, combined with goals and projected future state. Good situation awareness improves decision making in dynamic environments, and the way in which information is presented has a significant influence on situational awareness.²⁰ This is an exciting area of research and innovation, and the hope is to ultimately combine the medical knowledge of the CPMs with situational awareness and robust multifaceted context.

Measuring the impact of knowledge management is one of the most important and most challenging aspects of the program. Utilization data provides insights into the makeup of the user base and the types of information they most frequently seek. However, utilization metrics are insufficient to measure the real impact of knowledge delivery. A formal research program has been launched, and 2 studies are underway. One measures the effectiveness of the CPMs in standardizing practice, and the other measures the effect on physician behavior of delivering care recommendations through MEA. Through a partnership with Mayo Clinic's Center for the Science of Health Care Delivery, data are gathered and analyzed to provide a continuous improvement loop for the development of new knowledge and more effective delivery of knowledge to improve patient care. Specifically, the MEA prototype includes a mechanism to query EHR data and to measure and analyze practice variation. This process provides information that will allow continuous refinement of the CPMs and monitor progress toward practice standardization.

8 | CONCLUSION

A point-of-care knowledge resource developed to support an individualized approach to patient care has grown into a formal knowledge management program. This has been a key strategic initiative to focus the best of Mayo Clinic's multispecialty, multidisciplinary knowledge around the needs of the individual patient. Translation

of the textual knowledge into machine executable knowledge will allow integration of the knowledge with specific patient data and truly serve as a colleague and mentor for the physicians taking care of the patient.

REFERENCES

1. Glasziou P, Haynes B. The paths from research to improved health outcomes. *ACP J Club*. Mar-Apr 2005;142(2):A8-10.
2. Davis D, Evans M, Jadad A, et al. The case for knowledge translation: shortening the journey from evidence to effect. *BMJ*. Jul 5 2003;327(7405):33-35.
3. Lloyd FJ, Reyna VF. Clinical gist and medical education: connecting the dots. *JAMA*. Sep 23 2009;302(12):1332-1333.
4. Cook DA, Sorensen KJ, Nishimura RA, Ommen SR, Lloyd FJ. A comprehensive information technology system to support physician learning at the point of care. *Acad Med*. Jan 2015;90(1):33-39.
5. Michie S, Johnston M. Changing clinical behaviour by making guidelines specific. *BMJ*. 2004-02-05 22:50:47 2004;328(7435):343-345.
6. CMS. Stage 2 eligible professional meaningful use core measures, Measure 13 of 17. 2012. https://www.cms.gov/Regulations-and-Guidance/Legislation/EHRIncentivePrograms/downloads/Stage2_EPCore_13_PatientSpecificEdRes.pdf. Accessed July 21, 2016.
7. Cimino JJ, Jing X, Del Fiol G. Meeting the electronic health record "meaningful use" criterion for the HL7 infobutton standard using OpenInfobutton and the Librarian Infobutton Tailoring Environment (LITE). *AMIA Annu Symp Proc*. 2012;2012:112-120.
8. Peleg M, Tu S, Bury J, et al. Comparing computer-interpretable guideline models: a case-study approach. *J Am Med Inform Assoc*. Jan-Feb 2003;10(1):52-68.
9. Field MJ, Lohr KN, Institute of Medicine (U.S.). Committee to Advise the Public Health Service on Clinical Practice Guidelines, United States. Department of Health and Human Services. *Clinical practice guidelines : directions for a new program*. Washington, D.C.: National Academy Press; 1990.
10. Wang D, Peleg M, Tu SW, et al. Representation primitives, process models and patient data in computer-interpretable clinical practice guidelines: a literature review of guideline representation models. *Int J Med Inform*. Dec 18 2002;68(1-3):59-70.
11. Peleg M, Boxwala AA. An introduction to GLIF. *HL7 Winter Working Group Meeting*. Orlando; 2001.
12. Greenes R. Guideline Modeling. In: Greenes R, ed. *BMI 616: Clinical Decision Support*. Vol Week 4, Module 4 Tempe, AZ: Arizona State University; 2016.
13. Ohno-Machado L, Gennari JH, Murphy SN, et al. The guideline interchange format: a model for representing guidelines. *J Am Med Inform Assoc*. Jul-Aug 1998;5(4):357-372.
14. Wang D, Peleg M, Tu SW, et al. Design and implementation of the GLIF3 guideline execution engine. *J Biomed Inform*. Oct 2004;37(5):305-318.
15. Kawamoto K, Lobach DF. Design, implementation, use, and preliminary evaluation of SEBASTIAN, a standards-based Web service for clinical decision support. *AMIA Annu Symp Proc*.; 2005:380-384.
16. Tu SW, Campbell JR, Glasgow J, et al. The SAGE Guideline Model: achievements and overview. *J Am Med Inform Assoc*. Sep-Oct 2007;14(5):589-598.
17. Peleg M. Computer-interpretable clinical guidelines: a methodological review. *J Biomed Inform*. Aug 2013;46(4):744-763.
18. Quaglini S, Stefanelli M, Cavallini A, Miceli G, Fassino C, Mossa C. Guideline-based careflow systems. *Artif Intell Med*. 2000;20(1):5-22.
19. Patel VLK, David R. Cognitive Science and Biomedical Informatics. In: Shortliffe EHC, James J, eds. *Biomedical Informatics*. London: Springer-Verlag; 2014.
20. Endsley M. Toward a theory of situation awareness in dynamic systems. *Hum Factors*. 1995;37(1):32-64.

How to cite this article: Shellum JL, Nishimura RA, Milliner DS, Harper CM, Jr, Noseworthy JH. Knowledge management in the era of digital medicine: A programmatic approach to optimize patient care in an academic medical center. *Learn Health Sys*. 2017;1:e10022. <https://doi.org/10.1002/lrh2.10022>