

Using Clinical Decision Support to Make Informed Patient Care Decisions

Sep 19, 2008

Good afternoon, everybody, and welcome to this, the first in a series of teleconferences on clinical decision support. Today we'll be covering using clinical decision support to make informed patient care decisions. I am Teresa Zayas Caban, and I will be moderating. We have an excellent panel of speakers this afternoon, and before I introduce them, I would like to direct to you the bottom right of your screen, where a survey will pop up. We are using the survey to collect feedback on your user experience. Please fill that out.

Our first speaker this afternoon is Dr. Jon White who directs the health information technology portfolio at the Agency for Healthcare Research and Quality. Dr. White is responsible for setting the program for health IT projects and leads the HIT team at the agency. A board certified family physician, Dr. White received his medical degree from the University of Virginia and trained as a resident at Lancaster General hospital in Pennsylvania, where he received a national AAFP award. Today, he will be talking about the clinical decision support demonstration. Our second speaker is Dr. Rick Shiffman, professor of pediatrics at Yale University School of Medicine and associate director of medical informatics. His research involves development of tools and techniques that facilitate translation of knowledge contained in clinical practice guidelines into computer based decision support. He serves as the project director for the Glides project and will be presenting on replicable approaches to the development of ambulatory decision support. Our third speaker is Dr. Blackford Middleton, director of clinical informatics Research & Development at Partners Healthcare and Chairman of the Center for Technology Leadership. In this role, he leads the research and development group responsible for enterprise product development for the partners EMR and patient portal, and the enterprise clinical informatics infrastructure group. At CITL, he directs the value-based technology assessment research program and is director for the LMN fellowship program at Partners. He has been building and evaluating clinical decision support systems for all of his twenty plus year career and will be presenting today on the clinical decision support consortium project. Jon, I will turn it over to you.

I want to thank you very much for taking time out of your day to come listen to some very exciting projects, very interesting projects. Teresa told you about me, not going to repeat any of that. I do work at the Agency for Healthcare Research and Quality. We are part of the Department of Health and Human Services, and we're a small agency but with a big mission to improve the quality, safety, efficiency and effectiveness of healthcare for all Americans. As for myself and Teresa in particular, we work with the health IT portfolio at the agency. Our mission is to improve the quality of healthcare in the U.S. through better use of information technology. So that's the agency and what we're about. Today we're going to talk about our clinical decision support demonstrations, and the reason we're talking about this is because here in the U.S., we have great people working in the healthcare system, and great organizations in the healthcare system, and we also have a lot of great opportunities to improve the quality of care that gets delivered in our country for reasons that I am sure many of you have heard about. I am not going to repeat things like the article that says that 55% of the time people get the care they're supposed to get, and I am going to simply say that we really need to make the right thing to do the easy thing to do, which is something that my boss, Dr. Clancy, the director of the agency says often and says correctly. There are a lot of different ways to do that. In particular it is worth noting that clinical decision support has been applied to increase quality in patient safety, improved adherence to guidelines for prevention and treatment of illness, and to avoid medication errors in particular. The evidence shows that CDS can be used for a variety of purposes to improve the quality of healthcare. The meta-analysis of those reviews would indicate that that has happened well at a few select institutions but has not happened in a broad or widespread way across the country. So we know it is a good thing to do, but we need to know how to do it a little more broadly.

When we talk about clinical decision support, there are a number of good definitions out there. CDS can be defined (Shortliffe, 2006) as a computer based decision system that assists physicians in making decisions about patient care. Another definition comes to us from Wikipedia, but it is a rather good definition by Dr. Robert Hayward of the Center for Health Evidence which says clinical decision support systems link health observations with health knowledge to influence health choices by clinicians for improved healthcare. And a third potential defining set of words on the topic comes to us from the AMIA clinical decision support road map authored in part by Blackford Middleton and a number of other outstanding members of the community,

which says that clinical decision support provides clinicians, staff, patients or other individuals with knowledge and person-specific information, intelligently filtered or presented at appropriate times to enhance health and healthcare. Without getting into the details, that can encompass a variety of different things listed out here on the slide. So there are a number of ways to define it, but they all gather the same essence of what decision support is, so why isn't everybody using it? There are a number of barriers and those of you who you have worked with the systems in the past are going to be familiar with these. There is limited implementation of electronic medical records or computer provider order entry, difficulty developing clinical practice guidelines, lack of standards, poor support for CDS in commercial EHRs and challenges in integrating CDS into the clinical workflow. Often, underlying this is limited understanding of the people issues and organizational issues that go into making a successful implementation. So how do we get past that? Well, this is part of our job. Our job at the agency is to help create the knowledge, and synthesize it and disseminate it to the right folks so good things can happen.

In particular, around clinical decision support, we at the agency wanted to facilitate the development, adoption, implementation and evaluation of best practices using clinical decision support and to further enhance the nation's efforts to make evidence-based clinical knowledge more readily available to healthcare providers. We have great resources in things like guidelines.gov, the national guidelines clearing house and have been contacted by providers to say this is a great resource, but I need to figure out how to get it into my daily practice on a regular basis.

So we solicited for and awarded in early 2008 two demonstration projects. The objective of these projects was to develop, implement and evaluate projects that advance the understanding of how best to incorporate clinical decision support into healthcare delivery. The overall goal is to explore how the translation of clinical knowledge into CDS can be routinized in practices and taken to scale in order to improve the quality of healthcare delivery in the U.S., and all of this for the low, low price of \$1.25 million per project per year for up to five years.

Now, those of you who have done this work before will recognize that although as much as we want to be, we are not going to be able to be all things to all people for that amount of funding, so instead we set forward a number of key goals in 9 solicitations. We asked offerors to tell us how they were going to incorporate clinical decision support into EHRs that were certified by the CCHIT right now. CCHIT is the best mechanism at hand to standardize the health IT tools available to clinicians making daily decisions about healthcare.

We asked them to do that. We asked them to demonstrate cross-platform utility, and not just limit themselves to one certified product. We asked them to establish lessons learned from what they were doing for CDS implementation across the vendor communities which is a really important stake holder group to making this work, and to assess potential benefits and drawbacks to CDS. Also, we asked them to evaluate methods for creating, storing and replicating elements across multiple clinical sites and ambulatory practices—a topic that has been around for a fairly long time but that we felt was important to include in this work, and finally, for this particular set of work to translate clinical guidelines and outcomes related to preventive healthcare and treatment of patients with chronic illnesses. We wanted to try to get some breadth in the scope of conditions being addressed by these.

With that, that's the intent of what the agency wanted with the solicitation of these projects, and I would say that we went through a very rigorous selection process. We had a number of great applicants who did really good jobs. I wish I could have funded more. I think you are going to be really impressed to hear about where these projects are now, so with that, I will get out of the way and happily turn it over to Dr. Richard Shiffman.

Thank you, Jon. I would like to also welcome you as did Teresa and Jon earlier and thank you for taking part in this on a late Friday afternoon. The title of this presentation is A Systematic and Replicable Approach to Development of Ambulatory Decision Support. I am happy to introduce you to the GLIDES project. GLIDES is an acronym for guidelines into decision support, and it is a collaboration of Yale New Haven Health, Yale School of Medicine and the Nemours group whom I will have a little bit more to say about in a few slides.

This afternoon I am going to briefly describe goals and how we addressed the specific aims that Jon put forth in the RFP, and then talk in more detail about knowledge transformation and how we at GLIDES have addressed the issue of moving from guidelines into decision support. We began by defining our clinical objectives, then used a markup process involving the guideline elements model or GEM. From that we moved through XSL transforms. I will say a few words about action types and give you a preview of the user interface we're developing.

Very high level view of the GLIDES collaboration shows that we have a steering group made up of representatives from Yale and from Nemours that oversees the work of a guideline transformation group, a bunch of techies and clinical experts and implementation groups physically located both at Yale and at Nemours. We also have an evaluation group because to do all of this work and not understand how and why we got wherever we get, would be a real waste of the tremendous resources that AHRQ has made available to us.

The next slide would be a picture of the hospital, and a reminder that this is an almost 1,000 bed tertiary care hospital, including the children's hospital and the primary care center. It is a major teaching affiliate of Yale School of Medicine. Our pediatric primary care center provides care for about 8,000 inner city kids and about 28,000 visits annually. Nemours is a multi-specialty pediatric healthcare system consisting of more than 400 MDs and 4,100 staff in Wilmington, Delaware, Pennsylvania, New Jersey, and in Florida with bases at Orlando, Jacksonville, and Pensacola. In 2006 they had almost a million patient encounters taking care of a quarter of a million kids.

So, the specific aim to the GLIDES project, number one, is to implement evidence-based guideline recommendations that will address prevention of pediatric obesity and chronic management of asthma. Number two, to apply GEM, the guideline elements model and its associated tools to systematically and replicably transform the knowledge contained in these guidelines into a computable format. We want to evaluate the fulfillment of the goals and the effectiveness of the decision support tools in improving the quality of healthcare. The last and major specific aim is to disseminate what we learn and this activity is an important part of that.

Our project timeline overview looks at the two years for which we have been funded. We were really very close to schedule, having begun with project planning and knowledge transformation for both asthma and obesity guidelines. This is what I am going to be talking about mostly this afternoon, and we're well into the implementation process in the Yale specialty clinics. We will proceed to implement guidelines in obesity and asthma clinics at Yale and Nemours in Phase 2 and do primary care both at Yale and Nemours in Phase 3.

There is a great challenge in representing guideline knowledge electronically. Moving from the published guideline to a computer-based guideline implementation often involves a need for dually-trained, both medical experts and IT experts, who can do this translation project. The slide you're looking at shows the results of a study did by Patel in 1998 where she looked at collaborators from Stanford, Harvard and Columbia who were given a task in which knowledge engineers, the dually-trained folks, individually encoded guidelines for vaccine administration and workup of breast masses. They tested them by submitting standardized patients and found that different recommendations would be given for the same standardized patient.

So there is a black box that happens, which occurs between the published guideline and the computer-based guideline implementation, and one of our goals is sincerely to open that box and see what's inside to see if we can't make it systematically and replicably translatable. We do that using a four-part stack I will describe in a bit of detail, and you'll be able to compare and contrast our stack to Dr. Middleton's from whom we borrowed this idea.

There are a number of clinical objectives we identified for this project. It is very important to ground such a project in meeting clinical objectives and not just RFP objectives. Osheroff and Sittig in a publication from 2005 categorized the kinds of clinical objectives that decision support might be good for: for preventing errors both of commission and omission, for optimizing decision making, and for improving care

processes including documentation, patient education and empowerment, patient satisfaction and improving communication among caregivers.

So we convened conference calls to define our clinical objectives and define three criteria we would apply in order to select those objectives. Were these clinical objectives in fact addressed by the guidelines that we had selected, could they be facilitated by information technology and were they valuable? This slide shows a small chunk of the clinical guidelines we had looked at and where we next went to find recommendations within the guidelines that directly addressed these objectives.

This is here to remind you that there are established criteria for what constitutes a good quality guideline. One of those we helped to develop and published in 2003, and it is the conference on guideline standardization (COGS) checklist. It is a set of 18 criteria that can be applied to make sure the guidelines we're going to use have met some basic quality criteria.

Another is GLIA, guideline implementability appraisal. It helps to identify obstacles to implementation. Separate from guideline quality, it can provide feedback to guideline authors to help them anticipate and address obstacles before they release a draft guideline. It can also be used to help implementers in guideline selection and targeting attention toward anticipated obstacles. GLIA is available as is the COGS checklist on our website which you see here and will be broadcast again at the end of the presentation.

We faced a number of challenges. The NHLBI's asthma guideline update in 2007 is massive, more than 450 pages. NHLBI's effort at recording evidence quality and recommendation strength was commendable, but not uniform. There were multiple redundancies. Editing was irregular. We found some level of ambiguity, for example, a lot of recommendations were headed for children 0 to 4. Did they mean children who were really 0, that is at conception or birth, and does 4 include children who were 4 and 11/12 or only 3 and 11/12. Some of the choices that were made in the guideline recommendations were neither mutually exclusive nor exhaustive, nor well-defined. For example, they talk about interference with normal activity which could be none, minor, some, or extremely. The pediatric obesity guideline that came from the collaboration of AMA, HRSA, CDC and others had major deficiencies not the least of which was the absence of recommendation strength.

In moving from these guidelines, from a narrative to a semi-structured process which resulted in an XML file and a set of quality and implementability appraisals, we applied a tool called GEM cutter.

Gem cutter was developed following some work that we had done in the late 90s where we applied highlighters to guideline recommendations to parse them into replicable pieces. So, for example, in this recommendation we applied yellow highlighter to the decision variables, green highlighter to the recommended action, blue highlighter to the reason for that action, and purple highlighter to the strength of evidence. The trouble was there is so much information in guidelines we rapidly ran out of highlighter colors.

Along came XML which took us from a limited number of discreet colors to a virtually unlimited pallet.

XML is a multi-platform, web-based, open standard where we can create our own tags to enclose and describe text. For example, we can define a tag called inclusion criterion and put it around a term from the guideline like hematuria. What results is a human readable file that can be processed by a machine, and the markup activity can be performed by nonprogrammers, non-IT geeks.

GEM is the guideline elements model, a knowledge model for guideline documents. It was adopted first as a standard by ASTM in 2002. We updated it and changed it and restandardized it in 2006. GEM is intended to model the heterogeneous kinds of information that are contained in guidelines, and it is a multi-level hierarchy with well over 100 elements.

GEM Cutter II is the tool we use to parse guideline text into components of GEM. Call that process Gemifying. We create XML files and these XML files as well as GEM cutter can be found on our website (<http://gem.med.yale.edu>). The next step takes the semi-structured Gemified file and passes it through what

we call extractor transforms to create a semi-formal representation. This semi-formal representation includes statement logic: if-then statements with coded decision variables and coded action typing applied.

Extractor takes, for example, decision variables from their context in the guideline and presents them in a list. This affords an opportunity to judge their vagueness, their under-specification and decidability. It provides a comprehensive list of the trigger items that are going to be necessary for decision support activities and provides a measurable starting point for evaluation.

Here, for example, is a list of some of the decision variables from the asthma guideline. You will see 0 to 4 years of age, parental history of asthma, evidence of sensitization to foods. One of the things we have found quite useful is categorizing action types into a classification system we call action-types. It turns out that guideline authors don't have an unlimited pallet of things they ask us to do. They tell us to test, to monitor, to conclude, prescribe, perform procedures, to refer or consult, to educate our patients, document, dispose of our patients by admitting them, discharging or transferring, preparing our healthcare facility or advocating on their behalf.

The process for describing action-type patterns allows us to think about these actions in a systematic and replicable way, so, for example, any time there is a prescribed action called for in a guideline, there may be a need for drug information, for safety alerts, for formulary checking, dosage calculation, pharmacy transmission, patient education, and corollary orders. In the process of defining a decision support system around a prescribed recommendation, these are the kinds of things that may want to be included.

Finally, we move from a semi-formal representation to a formal representation in code, the code of the local electronic health record scripting language and a user interface that is designed to address local needs.

We see there being a knowledge pipeline that proceeds from knowledge in the universe through structuring that knowledge which is done by guideline authors and by guideline development system implementers through a zone of localization where local workflow is incorporated, and ultimately the knowledge is incorporated into a local electronic health record. That zone of localization is something we're currently exploring. We don't know how far you can push into that zone before you have to have pulling from the local facility at which decision support is to be installed.

Decision support can be delivered in a number of modalities, proceeding from the more static to the more dynamic. Simply providing documentation templates or prompts really does help to deliver a lot of decision support offering relevant data presentations. For example, display of relevant labs when ordering, order creation facilitators, providing reference information through an info button, providing reminders about appropriate care, and most dynamically alerts about drug allergies or interactions, and critical test result notification. I am going to show you a few screens that are currently being finalized by which we plan to apply these decision support building blocks in our decision support system.

So I suspect you may not be able to read this, but it is not hard to see that these are documentation templates that can be used to collect information about an individual asthmatic patient while you are in the clinic with that patient.

The asthma decision support that we're developing offers relevant data presentations—in this case, the step system for defining appropriate pharmacologic interventions for patients based on their severity of asthma classification and their level of control.

The system offers an alert in real time about its understanding of severity classification and an individual patient's impairment and level of risk. It offers an opportunity again to document the provider's classification and current understanding of the patient's level of control. So it is certainly possible for the provider to override the information provided by the system.

This is an order prescribing system, a facilitator of ordering, that makes it possible, simply by clicking the button that says order on the right, to order the appropriate kinds of medication that are defined by the provider. Thanks very much for your attention.

I will jump right in. Good afternoon, everybody. Thank you, Jon, and Rick for such a great setup. Thanks to the AHRQ for the funding, and to Jon White's brilliance in picking two projects which actually fit so well together, and thanks to Rick for a lot of great pioneering work in the GEM modeling efforts and you'll see how these projects I think will fit nicely together and hopefully help us advance the field.

There we go. Jon White has already outlined the objectives of the CDS demonstration project, so I won't belabor these points again. Jon outlined I think well the motivation and prior work and barriers that we've all experienced in this country trying to implement CDS, clinical decision support. One of my fundamental observations is that while we believe in the promise of HIT in EMR or CPOE or other electronic applications, it is very difficult to actually achieve the promise or achieve the value of HIT, and I think there are two fundamental reasons. One is that studies at the center for information technology leadership would suggest that in healthcare, he who pays for HIT is not he who gains. That's a separate conversation for another day, but we find that provider groups purchasing HIT only experience about 11% of the benefit and 89% of the benefit goes to other stakeholders in the health care ecosystem--payers principally among them. But another important aspect of clinical decision support, as Rick already put his finger on, is this problem of translating knowledge to clinical practice guidelines and then to clinical decision support implemented in HIT and not only implemented but effectively used. The CDS consortium has been formed to try to address many of these issues, and let me first highlight the members in the CDS consortium, and that's the clinical decision support consortium. We're fortunate to have enlisted the collaboration of the institute headed by Mark and his team, the Veterans Health Administration at Indianapolis, Indiana, with Brad, and Jason, the Kaiser Permanente Center for Health Research where Dean was, vendor partners, GE Healthcare along with NextGen, and OHSU and the University of Texas. Our primary goal in the CDS consortium is to attempt to assess, define, demonstrate and evaluate best practices for knowledge management and clinical decision support in healthcare information technology at scale, across multiple care settings, and across multiple EHR technology platforms. Our research objectives are schematically depicted here.

If you look at the box below, the knowledge management life cycle is an overarching goal for us to study and understand. We wish to extend and build upon Rick's great work with GEM to develop a knowledge specification which I will go into in a little bit more detail, and also to create a national knowledge portal and repository wherein members of the consortium can experiment with collaborative knowledge engineering across multiple sites of care and technology platforms, but further to access this knowledge in a ready-made forum at a variety of levels with specification that allows them hopefully to implement it much more easily than they can do currently and in current practices. To make that even as easy as possible, we hope to build upon the work of David and others at Duke and build web services, publicly available web services, so that if the technology at your site is ready to accommodate, could be inserted to provide decision support services from afar with assurance that the knowledge is kept up to date, is validated, and works well for the purpose of clinical decision support. We also will aim to evaluate each of our steps and objectives along the way and of course talk until we're blue in the face to disseminate our findings and the learnings we obtain.

The Office of the National Coordinator of Health IT in 2008 stated in its strategic plan that incorporating EHR functionalities and providing clinical decision support at the point of care is a key objective and that by 2010 certified EHRs should include clinical decision support. I think for us to achieve value with HIT investment and really move the needle on improving the quality of care delivered to this country with HIT, we have to address this knowledge management and decision support problem. Our focus areas will include diabetes, CAD, and hypertension screening and probably some other preventive care services as well.

The next two slides are fairly busy, but let me describe the teams that we put up now in this project to try to address our multiple objectives. First we have a Knowledge Management Lifecycle Assessment Team, which will look at the knowledge management and clinical decision support practices across the country in the member consortium sites. We'll try to derive best practices for knowledge management and clinical decision support in ambulatory care. The next team, the Knowledge Translation and Specification Team will help us achieve the four-layered representation Rick alluded to, and I am very happy to share with the Glide

project in a way that, through the portal, makes the best evidence and clinical decision support available to the widest audience possible.

The next team is the Knowledge Management Portal and Repository Team, which will stand up a national knowledge repository for the consortium that allows member sites in the consortium to access these artifacts at the various layers of specification in the knowledge framework and use them in their own clinical systems. Further, we'll do a collaborative knowledge engineering experiment to see if we can improve the rate at which knowledge and practice guideline is translated into consensus, clinical statements or the elements of a guideline decision support statement, and then translate it further into web services as I will describe further. The fourth team will make recommendations to the CCHIT, the clinical practice guideline community; to developers and anyone else who is interested about what we think are the best ways to do this in a practical and applied way to improve decision support in current HIT. Other teams are on the next slide: the CDS Services Team will be responsible at Partners Healthcare to build the web services--publicly accessible, publicly subscribable web services, which we will first test in practice at Partners Healthcare. For those of you who don't know, Partners is the parent organization for several of the Harvard teaching hospitals including Massachusetts General, and the Brigham Women's Hospital and several others. The web services will first be tested here and then in our long range project plan we aim to test them also in the Veterans Administration Vista medical records, and with NextGen and GE. The demonstration teams will evaluate the feasibility and the use of accessing knowledge via the knowledge portal as well as knowledge insertion from any level in the stack to a local EMR system. Another aspect of our research is to look at feedback from the field about how decision support is working and why isn't it when it doesn't work? So we aim to build also a CDS dashboard which for the end-user will give him or her a view of their own compliance with CDA alerts or other types of advisories coming from electronic medical records and further more feedback to the repository use characteristics on CDS performance characteristics for the knowledge engineers to review and then tweak the fundamental knowledge within the repository.

The dissemination team will focus on trying to keep all this ever present in the academic literature as well as in the industry forums we hope to attend, and we had to also create after we wrote the proposal -- we described also a joint information modeling team which has been very helpful in looking across all of the different teams to address the information modeling needs for practical application of our decision support framework or the knowledge representation framework in EMR.

We start with the KM life cycle assessment and from that we inform the knowledge translation and specification process. Both of those teams will develop requirements and use cases or assist with the development of the CDS web services, and the execution services. The demonstration projects of course will build upon the web services and the KM portal and help inform the design and development of the CDS dashboard. The KM portal is basically used as a collaboration resource and the knowledge access point for all these teams. From the combined learnings, we'll make recommendations to the CCHIT, HITSPI, and hopefully the measures community as well. We'll perform evaluations across the board in qualitative and quantitative manners where appropriate and disseminate to the best of our ability.

This is really born out of the need to be practical in decision support representation to allow an end-user or site to access knowledge in the way that is most comfortable to them. We feel that the industry is not ready across the board for web services nor is the industry capable of taking all the relevant clinical practice guidelines and developing logic and knowledge for implementation at their local sites. I want to avoid reproducing knowledge management and knowledge engineering exercises at each and every site implementing EMR, and we want to move people up the scale from localization of practice guidelines toward this semi-structured and abstract representation and hopefully ultimately a machine executable model of knowledge in the local EMRs. We believe this process will simplify the stewardship of CDS in a national knowledge repository, make it readily accessible to anyone who is implementing and able to consume publicly available web services. But others may access the knowledge in these intermediate states of representation, like the semi structured or the abstract, because that's what their system is ready to accommodate or their knowledge engineers and implementers can use. These layers have increasing precision and executability, of course, when you approach the machine execution end of the spectrum, and increasing flexibility and adaptability without any specification at the lower end of the narrative guideline. In a nutshell Rick has already outlined what these layers mean. The narrative recommendation layer is the

guideline layer, in its text form. The semi-structured recommendation layer breaks this down and, in fact, we hope to follow Rick's pioneering work with the guideline elements model to represent the guidelines similarly in GEM elements and then to develop the abstract representation layer, which begins to address the localization issues and context issues, but further specifies the knowledge. Ultimately, we want to develop the web service executable layer, which we can provide to you in a secure manner to provide local decision support.

This four layer representation is further extended with the idea of a knowledge pack. What are the essential critical elements for each and every piece of decision support to have available for the implementer to understand and then use in their implementation of this? Whatever the knowledge object is, the knowledge representation layer has four components:

- the data standards for the terminology used and the data standards for the data used including control of medical terminology, including concept definitions
- the logic specification,
- the functional requirement --this is a statement of what the EMR has to be able to do to consume the knowledge or express the knowledge at the local level in the EMR whether it is an order, alert or reminder or report or template, et cetera,
- we hope to build a CDS dashboard which will measure the impact and feedback for the knowledge engineers so that so they understand how their logic or their knowledge is being used in practice or not.

Why the multi-layered representation? I think I have alluded to our motivation in this regard several times already, so I won't dwell on this here except to say that we believe that knowledge has to be accessed in a way that is consistent with where the local site is and where the EMR technology is as opposed to one-size-fits-all. We think the four layers will allow local sites to access knowledge in the manner that is convenient and practical for them.

Here is another depiction of the knowledge artifacts, by layer, on the semi-structured recommendation representation, abstract rule or order set representation and then executable rules if the system can accommodate them or order sets or other art artifacts as the case may be.

As we're developing it, I want to recognize that the Complete CDS Knowledge Specification is a work in progress but we've tested a lot of these ideas in Partners Healthcare where we have built a repository and used services in our own EMR. It indicates and facilitates a variety of implementation methods for HIT and here is an example of a simple piece of logic--if the patient's cretonne is elevated, then avoid metformin--, and how we might elaborate and specify data, logic, function and measure, toward making machine interpretable what the actual detailed representation requirements will be.

We've done some of this work in the last few years at Partners Healthcare. We built a Partner's enterprise-wide, knowledge portal under the leadership of Dr Tonya Hongsermeier and in this repository or portal, we now make available the rules and content from all the different clinical information systems across Partners. This is a work in progress as well, but we're getting nearly complete on having all of the rule from CPOE, the different systems, the alerts and reminders, and even things like template specifications, report specifications, and the like all available in this repository. Our simple goal with this here is to make the best knowledge available to any system across our heterogeneous environment-- similarly to what we hope to achieve with the CDSC.

The way we used rules at Partners also has been funded by research dollars from AHRQ. This is a picture of a smart form we developed in our own EMR environment here. The EMR at Partners Healthcare is home grown. In the longitudinal electronic medical record we built the smart form environment with three components, the smart view or data display, smart documentation, and the idea of smart assessment, orders and plan. These three environments are actually all web parts driven by services and compiled on the fly based upon a knowledge base that is separate from the form itself as it is rendered in the environment. So, under assessment, you can see the highlighted elements here which draw the user's attention, the physician's attention to what needs to be done, and these are generated by the rules engine acting on the

rules in the repository, and make recommendation about all aspects of the patient's care and in this case for CAD and diabetes mellitus.

The smart view can also select intelligently the data you need to review for the particular problem encounter type at hand. And, for example, identify problems like the blood pressure is rising or significantly different or changed from last visit. That would result in an alert occurring on the right-hand side under the assessment panel, and the user could act on that by changing medications or adjusting medicine as appropriate.

The blood pressure measure here you can see is above goal, the average over the last two visits 130 over 80. The goal is less than that, so several alternatives are presented, and under the alternative, further guidance on how to select an appropriate anti-hypertension plan or other agent. In addition, we support the workflow as best we can because we feel oftentimes decision support fails because it doesn't fit into the workflow. So, you can see the medication orders, lab orders, referrals and handouts all fit into this fairly convenient order panel for this kind of visit in ambulatory care. The physician can adjust the medicine, order labs, make referrals and print handouts all in one convenient area, and that area, what goes in there is defined by the rule base, the knowledge base underneath. We've done a variety of work to-date, accomplishments are to date are listed here. These slides will be posted of course. You can read through this in more detail, but we have begun the life cycle assessment work. We've been deep into the knowledge specification work, building on the GEM work that Rick has done. We have specified the KM portal and are beginning to train users in how to use the KM portal and repository in the consortium. We've begun to think about the generalizations and analyzed the KM lifecycle assessment data to make recommendations to CCHIT. We have begun to specify the services for our own demonstration with the LMR and the modeling working group has completed analysis and design of a patient data model with the relevant terminology and data entities. We used the CCR as the specification -- the standard for our data and exchange formalism between remote EMR, and the decision support services we will create.

Our timeline for the first two years looks like this. We'll complete the knowledge management life cycle assessment, KTS work and the knowledge portal work, build web services and do the demo in our own LMR before the end of the second year. Thank you for your attention.

Thank you for the presentation. I would like to now open the panel up for questions. Please use the chat feature on your screen, and be sure to send questions to all panelists.

First question addressed to all panelists, it is from somebody who works in Alberta, and they have an Epic ambulatory system; they're considering a complete enterprise system. They're wondering about the differences between what you've been presenting for your project and something like Zynx?

Blackford, are you most comfortable addressing that?

I would to take a crack at it, Jon. I think the Zynx resource and the services they provide are excellent. I think the fundamental difference is we wish to put a body of knowledge, if you will, into the public domain, so it wouldn't be a vendor product. On top of the CDS consortium work, things might provide specialty decision support services or knowledge services for local decision support, but the fundamental difference is we aim to make at least the component of what's in the repository truly available in the public domain.

Rick, do you have anything to add?

Not really. This is a research project and not a vendor product, so our focus is really on trying to figure out how we can disseminate the information, not as a product but as a set of tools that individual organizations might use themselves.

The only other thing I would add to it is that these are demonstration projects, and they're meant to generate recommendations, but they're not yet meant to set policies or large national infrastructure in place. They can give us clues about how that would work, but it is not a product yet. The idea of a national, freely available repository of information is alluring on the one hand, not so much if you are a vendor of that knowledge, right, and I don't think we've yet worked out where the kind of boundary lines lie between what should the

federal government be doing to overall improve the quality of healthcare for all Americans versus what can the private sector contribute to that process and where are those lines, so we don't have the answers to that yet. I am betting by the time we get to the end of these projects, we'll be closer to good answers.

Thanks. We're getting a ton of questions here. Blackford, this question is directed at you. The presenter liked the idea of publicly accessible web services, and he would like you to talk more about what input you envision and which standards and what output? How would this be encoded or create executable instructions to a medication manager for CPOE?

Good question and complicated answer. What we do in our own environment is to draw from the information infrastructure in Partners Healthcare the relevant clinical data we need to have access to to infer or make an inference and return to the execution environment or the rendered form in the EMR. So data is accessed, inference happens in a rules engine and results are returned and expressed or rendered via the form at the point of care. Web services are getting easier and easier to build and serve up. On the other hand, how we create a national web service, if you will, for a potassium rule and how it is localized and data exchanged is still part of our research.

Thank you. One of our participant's works with inpatient CDS, and one of the challenges they encounter is ensuring roles and alerts remain current. I realize that you, Blackford and Rick, you're just getting under way, but can you speak to that challenge specifically as guidelines get updated and how that gets reflected in CDS?

I can take a crack at it. It is a very important question, and it is clearly going to be a difficult issue to completely resolve. The way we address it is by being guideline-based, and we depend on a document being released by somebody who is paying attention to changes in the knowledge that we can then put through our four-level stack and turn into a set of rules. We believe, although I don't have any empirical evidence to that effect, that having our rules stated as chunked decision variables and actions might facilitate the knowledge maintenance task, but as I say, I don't have any demonstrated proof that that works. It is a critical question. It comes up all the time, and certainly a centralized approach like the one Blackford is proposing has a better chance of being able to make changes quickly and dramatically than our approach which revolves around hard-wired IT activities. On the other hand, our approach is possible now with the infrastructure that's in place at our collaborative institutions.

Thanks, Rick. We have a couple of other questions. One of them is what sort of barriers exist in getting people to use this support and what strategy would you suggest would be more helpful to get providers to use decision support?

Is this where you tell them about the big stick you keep in your offices?

I wish. I think it is a very good question, and it goes to a lot of the cultural and organizational issues which we're not talking about directly. I can speak to what we do at Partners Healthcare, though, and I think clinical decision support in many ways starts at the top. It has to be an enterprise priority for whatever reason you might choose to have it be a priority, whether it is pay for performance or quality improvement or patient safety, or what have you. Then there needs to be clinical buy in at the grassroots level across the board in a material way that makes sense, and that means the decision support has to work well, and has to fit into the workflow, it has to not obstruct activity or interrupt activity, and has to be useful from the end-user's point of view which means it saves time or money or improves patient care in a way that really matters to the end-user. That's a very simple knuckle headed point of view on the topic which of course many folks have written volumes about and still remains a challenge for our country.

Thanks, Blackford. Rick, did you have anything to add?

I will say DITTO. I think you expressed it well, Blackford. I just want to point out we're working from an RFP, and we recognize that imposing decision support is harder to do than developing decision support for folks who have asked for it, so one of our first activities was a series of teleconferences with local stakeholders to define what they saw as clinical objectives related to asthma and obesity prevention that we might

address with our decision support system. That was our effort to get the buy-in that Blackford has so articulately described.

The only other thing I would add, this is not a case of if you build it, and they will come. If you build it well, then they'll probably come. So I think Blackford and Rick have outlined some great strategies. Also, both observed that we stand on the shoulder of giants. A lot of folks have written and done a lot of great work on this through the years. Another aspect of building it well means having information that's useful to the people who are providing the care and the people who are getting the care. We were asked a question from the group listening, from a colleague who said the definitions you use addressed clinical information, but don't necessarily address things like financial information, and isn't that just as important? The answer is, oh, yes, of course. That is really important. In fact, that's something that's even less obvious to the clinical users right now. We all know where to go find guidelines if we need them. We just can't get at them quickly and easily in a way that does not hurt our workflow. We often would have no clue where to go for financial information although we recognize that it is important. In the interest of keeping Blackford and Rick and their teams focused on work to move forward, we purposely chose to limit this, not just to clinical information but to guidelines in particular, so there is a whole big ocean out there for us to boil, and we'll eventually have to boil it to make it work well.

Thanks. Blackford, this question is probably for you. There are two sorts of related questions. One is how will localization be achieved in web services, and another participant is wondering if you'll have a plan for sharing the web service model through something other than directly through EHRs, and maybe you can give more detail about what you were talking about with the web service model.

Sure. Localization at some level has to be represented in the knowledge specification formalism. We believe that there is a way to accommodate within reason the variability that will occur at the local level in the specification of the knowledge itself. So within reason, again, standardization is part of the goal here, so we're not going to accommodate probably any localization, but we'll accommodate localization within reason, and then the usual mechanisms will apply. That is, some data that will be obtained if you're using the web service, some transaction will occur, and some inference will be returned to the local environment. Remember also, though, that this localization problem we know we may not be able to solve in all dimensions with the web service. Therefore, accessing the knowledge in a repository at the semi-structured or the abstract layer may allow a local knowledge engineer to do that hard part if it cannot be done by the web service, so that's an important part of the question here. How is knowledge best taken up and what level of specification is it taken up at, and most expeditiously across the range of options that we're going to present? I guess the second part of the question was about how a web service be delivered? I don't know if I have any insights on that yet because I think we're, as Jon said, doing experimental and developmental work for demonstration purposes. It is not clear exactly what will be the future of connectivity or information exchange and whether or not that will include any part of these kinds of knowledge-based services. It could.

Thanks.

Related question. How do you envision (not necessarily a question for what you're doing right now), small physician practices could be reached with these kinds of services, technologies that you're both describing?

Rick, do you want to start?

Well, I can try. It is a copout, I am afraid. Our collaborators include a number of small physician practices through the Nemours collaboration located in the Delaware valley, but these are not typical small practices. These are practices that have membership in the organization and have access which is not necessarily available to small practices. It is an important question, going to be a difficult problem to resolve because of the expense of the systems when the organization or government isn't subsidizing them. I think you probably have a similar situation with Partners, VA and your other organizations. You're not really reaching out to the small practice, are you, Blackford?

That's fair, Rick. We do also have within Partners the small office environment that is in the community, and they are using the EMR that is already consuming web services, et cetera but the question goes to a

couple of the critical dimensions of the demonstration projects. The vision we have is that tooling your EMR with the relevant decision support for your environment should be as simple as downloading the appropriate tax forms to Turbo Tax. Certainly that mechanism works well, standardized, et cetera. We don't have anywhere near the capability now across the landscape of installed EMRs, but I think we need to move toward that vision if we're going to actually get at this problem and solve it in any meaningful way.

Thank you. One of our participants was wondering if there is any collaboration with the GELLO project.

Why don't I briefly take a stab at that, and if you all have anything to add, you're welcome to. GELLO, which stands for guidelines expression language, is a great resource out there. We didn't explicitly require our offerors to use that particular language. We wanted to give folks an opportunity to come in with the tools that they saw best, so we didn't explicitly require it in the solicitation. So with that, I would ask Blackford or Rick, if you have anything to add.

Well, I can just offer that we are looking at it very carefully, and evaluating it for these purposes. It addresses some of our needs in terms of the expression of guidelines, but it doesn't address others.

Okay. So unless there is something further from Rick, I think that's enough of an answer.

Thank you. One of our participants is wondering how much training and follow-up training you feel you will need to do or will need to be completed for end-users and they were also wondering if end-users are part of your evaluation processes?

Another great question. Luckily, all of our end-users are currently using the electronic health record systems that are in place from the Centricity and Epic vendors, so we don't have any absolutely computer naive members. We're not looking to implement a new electronic health record system. We're looking to implement new decision support on top of existing electronic health record systems. That certainly simplifies our task by a considerable amount. That said, it will be important that our users understand what we hope to be quite intuitive decision support, and we will be evaluating that both qualitatively and quantitatively-- quantitatively, our success in delivering decision support that requires minimal but some training.

Blackford, at the AHRQ annual conference last week, Tim mentioned some of the required training that happens at Partners. Are you able to speak to that in more detail?

Yep. The question is an excellent one. I think we are taking the same approach as Rick is in that we're installing upon -- we're experimenting upon-- an installed user base, so there is not new training per se for the new decision support that we're introducing because it is only a marginal and relatively small add-onto what's already there. To the broader question, though, about training of EMR and training for CDS, it is absolutely critical to have well-described what the end-user will engage with and find useful. It is a very challenging problem, of course, to get physician time. What we find we have to do is not only offer the usual at the elbow support upon system design and implementation, but further to reinforce training on a regular basis and basically annually so the users know what's going on with the EMR. What are the new features, the new decision support, what have you, and that's a considerable but sizable but important expense to make sure you have accommodated.

Thank you. Sort of related to the earlier evaluation question, one of the participants is wondering whether ongoing assessments need to be made in order to determine the efficacy of following prescribed guidelines versus rendering an alternative treatment for a particular patient type?

Another great question. Evaluation of guideline-prescribed activity sincerely is something that has been underdone, and that's partly because guidelines prescribe processes of care, and the outcomes that may occur as a result of those processes are temporarily and spatially often quite distinct from the offering of decision support to a provider. For example, asking a provider to classify asthma and therefore choose a set of pharmacologic interventions won't be expected very soon to have a long-term effect on patient outcomes, on their long-term lung growth, which is one of the things we're interested in in pediatrics, on their quality of life. It takes awhile for these things to act. So far, we've been working on it in a two-year time

horizon and we hope to have some information that will help to validate outcomes relevant to these guidelines, but I don't expect we're going to have it within the two-year timeframe.

Thanks, Rick. One of the participants is wondering about clinical decision support and knowledge management—will this help a provider automate the creation of an order set for CPOE? Can you talk a little bit more about the difference between -- well, the works you're doing in terms of incorporating guidelines and how that may or may not tie into CPOE systems?

Do you want to start, Rick?

The concept of order set is something that really has been most developed in the in-patient setting. We think of writing prescriptions in the outpatient setting, and we think of order sets in the in-patient setting. Our scope in this project has been limited to the outpatient setting. That said, we're very definitely interested in making sure that appropriate pharmacologic selections are made for our patients with asthma, and a good bit of our decision support design has been towards funneling the information we collect into accurate conclusions about appropriate pharmacologic intervention. Those are going to wind up as prescriptions rather than as sets of orders.

I agree with Rick's comments and observations. The only thing I would add is sometimes there is a critical piece of logic which has to associate an order set with a diagnosis, a condition, or a pattern of care if you will, whether it is a disease, profile, or set of labs or what have you. So, in a way, that logic I think could also be represented in the approaches we're describing even though the order set it is a fairly simple decision support intervention.

I wish these guys had more time to talk about all the different ways in which decision support can be provided because they're doing some neat and innovative things, both within the projects as well as other projects that they're doing, but sadly not enough time.

One of the participants is wondering about what tool or set of tools are being evaluated for the semi-structured and executable layers of the knowledge stack? I think she's referring to your graphic, Rick, but I think Blackford can offer insight here, too.

Let me just start because my final slide somehow got lost, and it did have the URL for our project. If you pick up your pencil slowly, I will let you know that it is gem.med.yale.edu, and you will find the tools at that location. Specific information regarding this project is the same, gem.med.yale.edu/glides. The tools we developed are all available for use by the public without specific licensing from our website.

Should I move onto another question?

Go ahead.

So the one participant noted that studies of paper guidelines have often noted that published guidelines can contradict each other. Will any of the centralized services address guideline conflicts in an automated manner?

Well, it is like saying which of your children do you love better? Blackford, I will let you speak more to this since you have talked about the repository. I will briefly say that in our other work similar to this; we have not made an attempt to establish a gold star for one guideline over another. One could easily envision that. I personally wouldn't want to be responsible for doing that. You could also envision a rating system with our wonderful web 2.0 happening out there, but I don't readily foresee automated adjudication between different guidelines that address the same topic without a lot more thought put into it, so Blackford, further thoughts?

I agree totally, Jon, and to the questioner, I would suggest what we're trying to do is not to automate that guideline resolution or even prioritization process per se but to make the consensus process much easier to do than it is today right now. We have established a collaborative web 2.0 like engineering environment,

where folks can simply get together and resolve the ambiguities of guidelines and arrive at consensus much faster than they could before, but it is still a human mediated process.

Let me just echo that sentiment. I think this is something humans need to do, but I remind you there are tools that have been applied that can help you do this in a structured, systematic and replicable way, and I will refer you to the agreed collaborations tool, and the COGS checklist as well as the GLIA tool.

And I am just going to offer one quick thought. I didn't say this at the beginning; I thought about it but decided not to. When folks bring up comments like that, they often do it with a not necessarily unjustified concern that this sort of project moves us closer to cook book medicine, where the computer is doing the thinking for us. My hope is, after absorbing the complexities of what Blackford and Rick have been struggling with, and frankly that a lot of us have been struggling with, I think my hope is that there is an appreciation that the goal is not to have machines think for us, it is to give us the information that we need to make your clinical decisions in a timely way with the most up-to-date information, recognizing the fact that sometimes there is not good evidence about a particular decision, okay, or that there is conflicting good evidence about a particular decision. That is a much deeper issue for us to struggle with in healthcare.

Thanks. I don't see any new questions, so I think we'll actually end a little bit early today.

All right. I will just take a brief second and say thank you very much to the folks who have been listening for the outstanding questions and your attention and thank you to Blackford and Rick. We continue to admire your work, and are grateful to be working with you and thank you very much to Teresa for moderating us, an unruly group, and the questions as well.

Thanks, Jon, and before I forget, there will be a poll popping up at the right-hand of your screen. Please fill that out to give us feedback on the conference and help us in planning for future conferences. I would also like to remind folks this is the first in a four part series of teleconferences on clinical decision support. Please stay tuned for more information. The second teleconference is planned for October 27th from 2:30 to 4 p.m., and it will be focusing on the impact of clinical decision support on workflow, and I know a few of our participants have concerns about integrating CDS into work and issues of training, so we'll have great speakers lined up on October 27th. Thank you all and don't forget to fill out the poll.