A Virtual Medical Record for Guideline-Based Decision Support

Peter D. Johnson MB BS[a], Samson W. Tu MS[b] ,
Mark A. Musen MD PhD[b], Ian Purves MB BS MRCGP MD[a]
[a] Sowerby Centre for Health Informatics at Newcastle, University of Newcastle, UK
[b] Stanford Medical Informatics, Stanford University School of Medicine
Stanford, CA 94305-5479, USA

A major obstacle in deploying computer-based clinical guidelines at the point of care is the variability of electronic medical records and the consequent need to adapt guideline modeling languages, guideline knowledge bases, and execution engines to idiosyncratic data models in the deployment environment. This paper reports an approach, developed jointly by researchers at Newcastle and Stanford, where guideline models are encoded assuming a uniform virtual electronic medical record and guideline-specific concept ontologies. For implementing a guideline-based decision-support system in multiple deployment environments, we created mapping knowledge bases to link terms in the concept ontology with the terminology used in the deployment systems. Mediation components use these mapping knowledge bases to map data in locally deployed medical record architectures to the virtual medical record. We discuss the possibility of using the HL7 Reference Information Model (RIM) as the basis for a standardized virtual medical record, showing how this approach also complies with the European pre-standard ENV13606 for electronic healthcare record communication.

INTRODUCTION

In recent years, we have seen an explosion of interest in clinical practice guidelines and protocols. The use of guidelines promises to reduce inter-practice variation, to promote evidence-based medicine, and to contain the cost of health care. Studies have shown that guidelines are most effective when they are used to inform patient-specific decisions at the point of care [1]. Such usage of guidelines requires that decision-support systems use patient data from information systems of local sites to generate contextually relevant recommendations. Current guideline implementations take different approaches to the problem of data retrieval and normalization. Some systems, such as the CARE system at Regenstrief [2], grew out of and are tightly integrated with specific local clinical information systems. They represent patient data in highly institution-specific ways and guidelines authored in such formalisms are not readily transferable. Medical Logic Modules (MLM) written in Arden Syntax are designed to be transferable, but they leave the problem of getting and converting patient data into terms used by MLMs in "curly braces" whose implementation is institution specific [3].

In recent years, the EON project at Stanford Medical Informatics and the PRODIGY project at Newcastle's Sowerby Centre for Health Informatics have been collaborating to develop decision-support systems for guideline-based care. Using Stanford's experience in developing temporal-data mediators and Newcastle's experience in implementing advice systems embedded in the clinical information systems of multiple vendors, we have been developing the approach of defining *virtual medical records* (vMR) for guideline-based decision support. A vMR supports (1) a structured data model for representing information related to individual patients, (2) domains for values of attributes in the data model, and (3) queries through which guideline decision-support system can test the states of the patient. Based on the vMR data model, we have defined languages for expressing logical criteria that encode patient conditions relevant for checking a patient's eligibility for guidelines and for making decisions on the management of patient problems. Additionally, the vMR allows guideline authors to encode clinical guidelines using a rich and well-defined model of patient data and allows system developers to specify the requirements for guideline-specific storage, such as logging of guideline-based clinical assessments, decisions, and goals.

In this paper, we first discuss issues involved in defining a vMR, and then describe alternative approaches that the EON and PRODIGY projects have taken to define and use vMRs. We show that these vMRs are consistent with the European pre-standard for communicating part or whole of an electronic healthcare record (EHCR), ENV13606 [4]. Finally, we indicate how a standardized vMR may be derived from the HL7 Reference Information Model (RIM), thus opening up the possibility of a standardized methodology for embedding guideline-based decision-support systems in alternative electronic medical records (EMRs).

VIRTUAL MEDICAL RECORD FOR MODELING GUIDELINES

An institution's real patient medical record is a complex legal entity containing detailed financial, demographic, and clinical information about the patient, together with attribution and audit trail data. In defining the vMR to provide decision support for guideline-based care, we do not attempt to create a data model that replicates everything that an EMR holds, but only those

distinctions necessary for modeling guidelines and protocols. That fewer distinctions are required makes the process one of simplification, thus making it more feasible. For example, an asthma guideline may require that we know whether the patient is currently taking an 'inhaled β_2 agonist' and at what dose. Other attributes of the prescription in the EMR (e.g., number of refills, identity of prescriber) are irrelevant.

In addition to defining the categories and the structure of patient information that can be queried, a vMR has a data dictionary that contains terms that can be referenced in modeling a guideline. These terms may come from a standard vocabulary such as ICD10 or NHS Clinical Terms, but the purpose of the term is to denote medical conditions or interventions being discussed. The previous example is typical in that the concept being queried, inhaled β_2 agonist, is more general than those usually present in a record (e.g., specific drugs such as albuterol). To satisfy this query, the guideline system has to look for all medicinal products that are β_2 agonists, and that have an intended route of inhalation. The necessary translation is done by a one-to-many mapping from the general concept to those that are more specific. Such a mapping is easier in so far as problems of marginal differences in the exact meaning of individual specific terms do not arise because of the generality of the query term. On the other hand, other mismatches between the data model and terminology of a vMR and those of a real EMR may make the mapping process an arduous one. For example, a concept used in a guideline (e.g., Helicobacter-positive gastric ulcer) may have no exact equivalent in any standard vocabulary. It may have to be defined as a combination of other concepts or it may have to be queried from a human user. There may be differences in units of measurements (e.g., patient height) or temporal granularities (e.g., abstracting an episode of hyperkalemia from several time-stamped observations). Results of tests done in different settings (e.g., outpatient vs. inpatient) may not be comparable. Any system that implements published clinical guidelines for an EMR has to solve such problems [5].

Finally, the interaction of a guideline system with the medical record is not necessarily a one-way process. In addition to querying the medical record for information on observations, problems, and procedures, there is a requirement for recording clinical assessments (e.g., hypertensive on maximum dose beta blocker, poorly controlled) and decisions made by the clinician based on the recommendations of a guideline (e.g., start second anti-hypertensive agent), as well as the actions that result (e.g., provide prescription for Diltiazem, order an investigation). These interpretations, decisions and acts need to be recorded through the virtual medical record.

EON SOLUTION

The patient-data model in EON defines the classes and attributes of patient information required by the rest of the system. We model patient data as static, time-stamped, or associated with a time interval during which the information represented by the patient data is valid. The model consists of a *Patient* class whose instances hold demographic information about specific patients, a *Qualitative_Entry* class that describes qualitative observations about patients, a *Numeric_Entry* class that represent results of quantitative measurements, an *Adverse_Event* class that models adverse reactions to specific substances, a *Condition* class that represents medical conditions that persist over time, and two intervention classes, *Medication* and *Procedure*, that model drugs and other medical procedures that have been recommended, authorized, or used.

Based on the patient data model, the EON guideline modeling system contains multiple criteria languages for expressing patterns of patient conditions [6]. For ease of guideline encoding, we have created a number of templates that allow easy specification of simple criteria such as *presence of bronchospasm within the last 3 weeks*. For querying complex temporal patterns that may involve abstractions from primitive data, we have built a general temporal query language [7]. For generalized criteria such as *presence of an authorized medication that is contraindicated by some medical condition*, we use an expressive constraint language based on the Knowledge Interchange Format.

For representing domain concepts (such as *bronchospasm)* used by guidelines and protocols, the EON system creates *medical-specialty models* consisting of taxonomic hierarchies of concepts and relationships among them. Domain concepts in these models may contain guideline-specific knowledge. For example, the concept *ACE Inhibitor* includes information about indications and contraindications that is described in a specific clinical practice guideline.

We have deployed the EON architecture in an advisory system for management of hypertension, known as ATHENA at the Veterans Affairs Palo Alto Health Care System [8]. We use several data mapping and transformation techniques in this application. First, to perform the one-to-many concept mapping described earlier, we augment the medical-specialty model so that appropriate ICD9 codes are represented as specializations of disease concepts in the model. To overcome data mismatches (such as unit and temporal granularity mismatches), we created a mapping knowledge base and associated software modules that transform the schema and terminologies of the data extracted from VA's VISTA system to a form consistent with the patient-data model assumed in EON [9].

PRODIGY SOLUTION

The model of patient data in PRODIGY [10] is similar to that of EON. Patient data are either time stamped or associated with a time interval. Demographic data are held in instances of a *Patient* class. A *Note_Entry* class is equivalent to EON's *Qualitative_Entry*, whereas an

Investigation_Result is equivalent to the *Numeric_Entry* class of EON. Rather than model *Adverse_Event*, the PRODIGY model has a *Sensitivity_Entry*, which records that a patient is known to be sensitive to an agent. The authorization for medication is modeled as a *Therapy_Entry*.

Because PRODIGY is used on more than one vendor's system, we defined a standard query interface [11]. The use of this interface requires the vendor to provide a mediation component, which maps their proprietary database structure into the above classes, and which exposes the standard programming interface. Because these systems use different terminologies (NHS Clinical Terms versions 1, 2 and 3 and a variety of proprietary drug terms), a mapping problem arises. To enable use of only one shared guideline knowledge base and execution engine, all queries are defined in PRODIGY Query Terms – these are domain-specific, high level terms, (e.g., 'diagnosis of renal failure') which map onto sets of the target clinical system vocabulary. The presence of one term in the set will satisfy the more general query term. A separate knowledge base of the mappings of these query terms to the relevant vendor system vocabularies is provided, so that queries are passed across the interface as PRODIGY Query terms, and resolved into the local vocabulary by the vendor.

Early users of PRODIGY requested that we record explicitly in the EMR goals agreed, decisions made, and clinical interpretations made in conjunction with the guideline, building up a story of guideline use over time. This requirement results in new classes of data to be written into the EMR: (1) clinical interpretations made by the guideline system or chosen by the user, e.g., 'hypertensive on maximal dose beta blocker,' and (2) a recording of guideline decisions made – these are usually intended actions, e.g., 'start second anti-hypertensive'. Again the term in the guideline is a PRODIGY term, so mapping is required to the target vocabulary. As these terms are being written to the record, there is a one to one relationship, and the terms are not generalized.

For testing the state of the patient, the PRODIGY provides a set of standard criteria templates, which authors populate using PRODIGY Query terms. These templates allow simple criteria to be defined such as *diagnosis of asthma in past 6 months*. Criteria based on abstractions from primitive data can be created, such as *10 year cardiovascular risk > 15%***.** However, complex temporal patterns cannot be defined.

TOWARD A STANDARD VIRTUAL MEDICAL RECORD

The success of the EON and PRODIGY projects in using a vMR for the purpose of authoring guidelines and in mapping the vMR to multiple vendor systems demonstrates the possibility of implementing a single guideline-based decision-support system in multiple clinical environments. Clearly this capability would be of major benefit, as the cost of developing guideline knowledge bases and execution systems is high. Standardizing the data model of the virtual medical record and creating a methodology for the mapping of diverse medical record systems to this standard would alleviate one of the obstacles in providing patient-specific guideline recommendations at the point of care.

Any effort to create a standard vMR for guideline-based decision support needs to make sure that it is consistent with EMR standards that may be widely adopted. In 1999, the European committee for standardization (CEN) approved a pre-standard, ENV13606, for communicating part or whole of an electronic healthcare record (EHCR) [4]. We now describe ENV13606 briefly and show that the EON and PRODIGY vMRs can be mapped into the ENV13606 structure. The result suggests that the EON and PRODIGY data models are generalizable, and that we can use the European pre-standard as a source for data elements that might be missing from the current EON and PRODIGY models.

CEN ELECTRONIC HEALTHCARE RECORD (EHCR) COMMUNICATION PRE-STANDARD

The CEN EHCR communication pre-standard defines an architecture for the EHCR. The intention is that the pre-standard will enable effective communication of healthcare record information through faithful preservation of its content and context. The fundamental entity in ENV13606 is the *elementary healthcare record entry* represented by an instance of a *data_item* subclass. These subclasses map onto vMR concepts:

- *structured coded data item;* An item of healthcare related information represented by a term from a vocabulary. Examples are symptoms, signs, problems, diagnoses and procedures.
- *quantifiable observation data item*; The result of a single clinical measurement or laboratory investigation.
- *medication data item;* A treatment, with a code for the drug, or other prescribable product.
- *person identifier data item***,** Demographic Information about a person, e.g., sex, age

Annotations are additional attributes to data items to specify, for example, the LIFE_CYCLE of a procedure: done, in-process, or planned. POTENTIALITY defines an observation as actual, a goal, or predicted. The KNOWING_MODE defines an observation as objective, physical examination, or subjective opinion.

Data items can be related by *link set items* to convey causality, motivation, interpretation and goal relationships.

Data items are aggregated in various ways. Data items from one place or time of care form a *Composition*, while various task specific views are subsequently created as *Headed Sections*, e.g., Active problems.

Possible vMR Concept	**EON**	**PRODIGY**	**ENV13606**	**Contains**:
Patient	Patient	Patient	person_identifier_data item	Demographic data
Qualitative Observation	Qualitative_Entry	Note_Entry	structure_coded_data_item	e.g.,symptoms, signs
Quantifiable Observation	Numeric_Entry	Investigation_Result	Quantifiable_observation data item	e.g., height 1.78m
Medication Authorization	Medication (authorized)	Therapy_Entry	Medication_data_item	e.g., Atenolol 100mg one a day, 28 days
Procedure (done)	Procedure	Note_Entry	Structure_coded_data_item Annot. LIFE_CYCLE = done	e.g., Pancreatectomy
Allergy State	Adverse_Event	Sensitivity_Entry	Structure_coded_data_item Archetype: allergy state	e.g., sensitive to penicillin
Clinical Assessment	Qualitative_Entry	Scenario_Entry	Structure_coded_data_item Annot: KNOWING_MODE = assessment or opinion by author	e.g., poorly controlled hypertension
Goal			Structure_coded_data_item or Quantifiable observation data Annot: POTENTIALITY = goal	e.g., aim to get BP less than 140/85
Planned Intervention	Procedure (ordered)	Act subtype	Structure coded data item Annot. LIFE_CYCLE = planned	
Decision		Action_Decision	Structure coded data item Archetype: decision activity	e.g., start 2nd line therapy

Table 1. Mapping of record types in virtual Medical Record

Table 1 shows how classes in the EON and PRODIGY vMR can be mapped into corresponding structures in the ENV13606 pre-standard, defining a set of ten classes which make up a possible set of standard vMR concepts. *Patient* holds demographic data, of which age and sex is most important. *Qualitative* and *Quantifiable Observations*, *Medication Authorization*, *Procedures* and *Allergy State* have been discussed in EON and PRODIGY, and map directly into ENV13606. There is less convergence in EON and PRODIGY on the recording of subjective clinical assessments, management decisions, goals and planned interventions, although the requirement for recording these has been clearly demonstrated. These concepts are well handled in ENV13606 so we propose to adopt this set in our work.

HL7 REFERENCE INFORMATION MODEL AS BASIS FOR A VIRTUAL MEDICAL RECORD

Even if the vMR we develop conforms to an EMR standard, it is unlikely in the near future that many local information systems will comply with the standard. We will always need to map data models of local information systems into that of the vMR. We believe that a path toward a methodology for mapping diverse EMRs to a standard vMR exists through Health Level 7's Reference Information Model (RIM) [12].

HL7 is an organization whose mission is to provide standards for the exchange, management and integration of clinical data. The HL7 RIM is an information model that is the source for the data content of all HL7 messages. The model is abstract, using a few, base concepts such as *Act*, *Entity*, *Role* and *Participation*, from which more specialized classes can be derived. *Entities* are used in roles – a *Patient* is a *Person* in a role of patient. *Entities* in roles can *Participate* in *Acts*. *Act* has many subclasses; such as *Observation, Procedure, Medication*. An *Act* has a *mood*, which allows us to say whether it is a definition, an intent, an order or an event. Thus *Act* allows us to record goals, planned procedures, as well as concrete observations [13].

As the RIM is a fundamental model; it is not directly usable as a representation of patient data relevant to the task of guideline-based decision support. Taking the *Medication* act as an example, this act does not itself define a medication order in the EMR sense. The *Medication* act, with a mood of 'order' can be used with the participation of a *Patient* (a *Person* in the role of patient) and a prescriber. It will have many *Act_relationships* e.g., one of type *component* to a *Patient_Encounter* act. Additionally, the RIM contains many classes (such as *Financial_act*) and attributes (such as *valuables_location_desc*) not relevant for clinical decision making.

HL7 provides a methodology for deriving domain-specific messages from the RIM. To create an HL7 message the required classes from the RIM are combined in a Refined Message Information Model (R-MIM), creating new relationships and imposing domain-specific constraints on attributes of the classes. This is a process of creating higher-level models for particular use cases. An example R-MIM for a

medication order is constructed from thirty two objects from the RIM.

The virtual medical record presented here can be seen as a set of R-MIMs, defined using the HL7 methodology. As yet, a standard set of R-MIMs for HL7 is not finalized. Examination of preliminary examples suggests that, as expected, those for the vMR will be simplified versions of standard R-MIMs. For example, for medication order, no new distinctions need to be added to the standard HL7 medication order R-MIM, but several parts of the model (e.g., the participation of a pharmacist in a verification function) are not required.

DISCUSSION

The heterogeneity of legacy schemata and local vocabularies is a major reason that software systems often are not interoperable in clinical settings. In PRODIGY and EON, we provide a logically consistent virtual medical record and concept model (1) that can be used for specifying guideline-related decision criteria and actions, and (2) that can be mapped to the information sources of the institution where we deploy decision-support systems. This approach provides a clearly defined process for integration with a pre-existing clinical information system and it enables reuse of the decision-support systems for guideline-based patient care. We have shown that the vMRs in EON and PRODIGY are consistent with the European CEN pre-standard for electronic medical records. This standard is complementary to the HL7 RIM, in that it takes a higher level, semantic view of the record, but does not specify the low level representation. All of the concepts required for decision support are present in this standard.

The HL7 Clinical Decision Support Technical Committee is undertaking the task of developing a virtual medical record that transforms the RIM into a view of patient data that can be utilized by decision-support systems. The community effort to develop a standardized virtual medical record is an exciting development. In this paper we have shown how we are taking preliminary steps to use the HL7 methodology to create views of medical records that supply the distinctions necessary for decision support. With the adoption of a patient data model conformant to the HL7 RIM, the problem of transforming legacy data models and schemata decomposes into problems of determining (1) how the legacy systems can export their data in HL7 message format and (2) how a guideline-based decision-support system can process patient data encoded as HL7 messages. This approach will allow for one-to-many mappings instead of the many-to-many mappings between decision-support systems and clinical information systems currently required.

Acknowledgments

The PRODIGY project is funded by NHS Executive Primary Care. The EON project has been supported in part by grant LM05708 from National Library of Medicine and a grant from FastTrack Systems, Inc.

REFERENCES

[1] Grimshaw JM, Russel IT. Effect of clinical Guidelines on Medical Practice: A Systematic Review of Rigorous Evaluations. *Lancet* 1993;324: pp. 1317-22.

[2] McDonald, CJ. Action Oriented Decisions in Ambulatory Care. Chicago: Yearbook Medical Publisher. 1984.

[3] Hripcsak G, Clayton PD, Pryor TA, Haug P, Wigertz OB, Van der Lei J. The Arden Syntax for Medical Logic Modules. *Proceedings of the 14th Annual Symposium on Computer Applications in Medical Care*. Washington, D.C., 1990: pp. 200-204.

[4] European committee for standardization (CEN) ENV13606:1-4 Electronic Healthcare Record Architecture, 2000.

[5] Rector AL, Johnson PD, Tu S, Wroe C, Rogers J. Interface of Inference Models with Concept and Medical Record Models. Submitted to AIME 2001.

[6] Tu SW, Musen, MA. Modeling Data and Knowledge in the EON Guideline Architecture. Submitted to MedInfo 2001.

[7] O'Connor MJ, Tu SW, Musen MA. Applying Temporal Joins to Clinical Databases. *Proc AMIA Symp* 1999; pp. 335-339.

[8] Goldstein MK, Hoffman BB, Coleman RW, Musen MA, Tu SW, Advani A, Shankar R, O'Connor M. Operationalizing Clinical Practice Guidelines While Taking Account of Changing Evidence: ATHENA, an Easily Modifiable Decision-Support System for Management of Hypertension in Primary Care. *Proc AMIA Symp* 2000: pp. 280-284.

[9] Advani A, Tu S, O'Connor M, Coleman R, Goldstein MK, Musen M. Integrating a Modern Knowledge-Based System Architecture with a Legacy VA Database: The ATHENA and EON Projects at Stanford. *Proc AMIA Symp* 1999: pp. 653-657.

[10] Johnson PD, Tu S, Booth N, Sugden B, Purves IN, Using Scenarios in Chronic Disease Management Guidelines for Primary Care. *Proc AMIA Symp* 2000; pp. 389-393.

[11] Johnson PD, Tu SW, Jones N. Achieving Reuse of Computable Guideline Systems. Submitted to MedInfo 2001.

[12] HL7, HL7 Data Model Development, (2001), http://www.hl7.org/library/data-model/.

[13] Schadow G, Russler DC, Mead CN, McDonald CJ, Integrating Medical Information and Knowledge in the HL7 RIM, *Proc AMIA Symp* 2000: pp. 764-768.