

The HL7 Reference Information Model Under Scrutiny

Gunther SCHADOW^{a,1}, Charles N. MEAD^b, D. Mead WALKER^c

^a Regenstrief Institute, Indiana University School of Informatics., Indianapolis, USA

^b Oracle, Redwood City, USA ^c Mead Walker Consulting, Downingtown, USA

Abstract. The Health Level 7 (HL7) Reference Information Model (RIM) was once incepted as an object oriented information model to harmonize the definition of HL7 messages across different application domains. On the heels of the hugely successful HL7 version 2, version 3 and the RIM has received significant attention and credit and in turn is increasingly subjected to criticism. In this paper the authors, who are among the chief designers of the RIM, respond to the major points that have been raised against the RIM in the published literature. We find that much of the criticism is based on misunderstandings and differences in point of view. We wish to advance the dialogue in the hope that when we account for those differences, effective critique may lead to real improvements of the standard. Keywords: Medical informatics, HL7, ontology, electronic health records.

1. Introduction

At the core of the HL7 version 3 standards development methodology is the Reference Information Model (RIM), which is a static object-oriented model in UML notation. The RIM serves as the source from which all specialized HL7 version 3 information models are derived and from which all HL7 data ultimately receives its meaning. This is to establish semantic interoperability across a vast and growing number of subject domains (e.g., laboratory, clinical health record data, problem- and goal-oriented care, public health, clinical research, etc.), which are loosely but critically related. The RIM was first conceived as a *data model*, where all data elements known from HL7 version 2 and some large electronic health record data models were put on a single information roadmap. In an iterative process of harmonization, analysis, unification and extension of scope, today's RIM emerged as an abstract model, which defines the grammar of a *language* for information in healthcare. As shown in Figure 1, all data is in a form in which Entities (e.g., people places and things, nouns) are related in Roles (relators) to other Entities, and through their Participations (prepositions) interact in Acts (verbs). Through ActRelationships, networks of structurally or logically related Acts are formed, expressing composition, reason, order-fulfillment, data-derivation, etc.. The RIM outlines the logical form of data, but it has specific extension points, at which users and other organizations can contribute content dynamically. This is done through domain-specific terminologies and using a data element definition framework which is built into the RIM directly. For example, all clinical data systems will have a "master file" which defines most clinical data elements, and the HL7 RIM reflects that. [1]

¹ Corresponding Author: Gunther Schadow, Regenstrief Institute, 1050 Wishard Blvd., Indianapolis, IN 46202, USA. gschadow@regenstrief.org

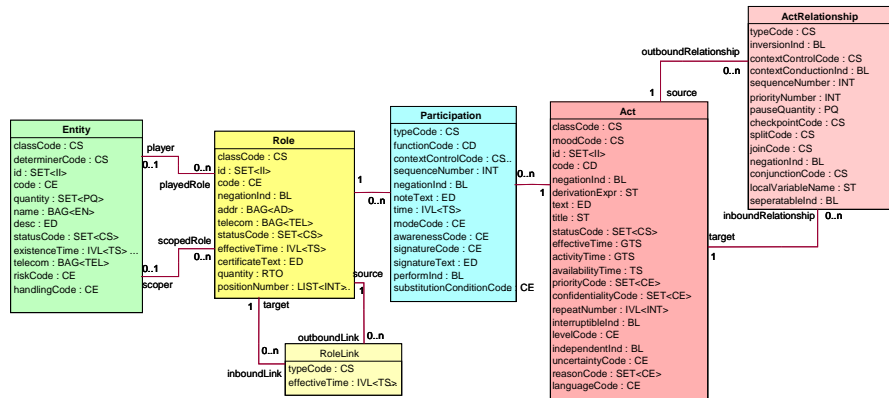


Figure 1: UML Class Diagram of the Reference Information Model “Backbone”

2. Technical Criticism

In surveying the literature (using PubMed, Google), we found that technical criticism is delivered only in passing, rarely in a fashion in which it could be dealt with to either improve the quality of HL7 or to dissipate the concerns. A paper by Fernandez and Sorgente is symptomatic; it claims that HL7 violates the Unified Modeling Language (UML), violates some ostensibly universally accepted object-oriented principles, and that it is not practically implementable. In truth all HL7 models are UML models and maintained as UML by the HL7 development tools. Domain-specific models may first be designed as unconstrained UML *domain analysis models* (DAM) and then transformed into a highly constrained normalized UML model where all classes in the constrained model must specialize one of the RIM classes and all features (attributes and associations) must derive from one of the RIM features. This guarantees that each instance-structure that complies with the specialized model also complies with the general model. This facilitates interoperability between interfaces for special domains on the one hand and general HL7 interfaces (such as for data warehouses) on the other, and thereby between two specialized interfaces where they overlap.

When mapping the special domain model to the RIM, we invoke an operation called *refinement* or sometimes “cloning”, where RIM classes are shown as separate specialized classes on a new page as if they were de-novo created classes (such graph-refinement or folding operations are well established in computer science.) Some specialized classes may then not use certain attributes or associations from its original RIM class, because, while these features logically apply, the domain specialists had no business case for them in their design. Some critics perceive this as “deletion” of inherited features (which is frowned upon). However, this is not deletion but *projection*, a well-established operation in relational database theory and practically used in virtually all integrative information systems: through projection one maps the needs of a special application to the general (enterprise) database. Specialization through constraints is fully compliant with the UML notion of specialization.

HL7 has used UML models from the beginning, but has later developed another visualization technique. These “block diagrams” are easily traceable to UML, but

visualize the standard RIM patterns and important constraints more concisely (e.g. vocabulary domain constraints). UML would use separate boxes with constraint annotations in the form of OCL expressions, which is much more verbose, provides less guidance, and is not suitable for the domain experts who are not computer scientists. HL7's deep experience with making model-based development work with domain specialists often creates the impression HL7 is "re-inventing the wheel" rather than using whatever is called the "industry standard" of the time. The truth, however, is that HL7 version 3 has been based on an object-oriented methodology since its inception in 1995/1996. HL7 not only actively adopted the best industry standard tools and methodologies available at the time, but also engaged the communities which developed these methodologies. When the HL7 version 3 project was launched, it had to pioneer a model-driven standard development methodology almost a decade before the first mentioning of the OMG's "Model Driven Architecture" (MDA). So, HL7 has not "re-invented" but rather originally invented many "wheels" to meet its requirements and vision at a time when the need for such wheels was not generally accepted.

That in the end the HL7 version 3 specifications may induce fear among some implementers is due to its mission of establishing semantic interoperability in loosely coupled systems. Thus, HL7 specifications address many of the practical difficulties with real world medical information processes (e.g., incomplete information, thing identification, duplicate record problems, uncertainty in data, etc.) directly, rather than deferring them to special applications. Even fundamental HL7 specifications such as the HL7 Version 3 Data Types do not remain silent about these ever challenging issues (such as uncertainty or incomplete information) where most general computing technology does not mention any of this. Implementers of special interfaces need not be concerned with the complexity that comes with HL7's full specification, as simple things usually are simple in HL7. But for more general HL7 interoperability, these issues are relevant, and need to be addressed (which would be more difficult without HL7.) Those who invest the appropriate effort in implementing general interfaces, such as the HL7 Java SIG or the Oracle Health Transaction Base, find that it can be done, that it works quite well, and they feed their improvements back to the working group.

3. Conceptual Design – Flaw or Virtue?

A range of conceptual doubts are raised, much of which have to do with the myth that HL7 is based on antiquated ideas of U.S. billing messaging, and that it therefore cannot deal effectively with ones favorite new challenge, be that "security" or the "electronic health record" (EHR). In truth HL7 has the most extensive collective experience of healthcare computing, specifically covering practical EHR; and security requirements are indeed considered very carefully. So, instead of arguing against myths, we shall focus in the rest of this paper on a core body of criticism raised by Smith et al. [2,3,4]

3.1. Incoherent Specification

Smith et al. seem to struggle with HL7's imperfect specification. Smith's ironic use of the term "exegesis" is not meant as a compliment. Smith rightfully scrutinizes the language of the HL7 definitions and discovers questions which are often overlooked in practice. Imperfections and inconsistencies in the HL7 documentation exist and need to be addressed, however, as much as one might try, inconsistency and ambiguity are

deeply unavoidable in constructing a collaborative volunteer-based standard, which brings together a wide range of people from different backgrounds. As different people edit parts of the specification, inconsistencies in form and quality may emerge; as some ambiguities are clarified, other previous systematic ideas may be corrupted; and well-meant glossary entries may cause confusion. Sometimes irreconcilably opposed conceptualizations may coexist and one resorts to vague or ambiguous language in the interest of moving forward in areas where parties can not agree. Thus, standards are akin to the body of laws in pluralistic societies, which at any given point in time seems imperfect, incoherent and even contradictory. Over the long term, however, the engagement of different viewpoints in the development of the standard yields the best result because it is practically relevant and there are no alternatives to such consensus.

3.2. “Double Standard” – Reference Information Model vs. Reference Ontology

Smith’s criticizes that the RIM blurs a distinction which should be made between *information model* and *reference ontology*. An information model defines what we record and communicate about the world, whereas the reference ontology would model the world itself. Smith shows how the definitions of the RIM elements switch between object-language (e.g., “person is a human being”) and meta-language (person-record representing information). In this, HL7 has followed common practice in object-oriented analysis casting the result of the real world analysis into information model designs, using classes that bear intuitive names, such as “Person” or “Procedure,” and yet being aware that no person will ever be stored in a computer system. Instead, HL7 documentation implicitly assumes that computer systems deal only with records of information about these real world phenomena. Since HL7 is about information management, some of its features (e.g. Entity statusCode) are about such management functions, but beyond that HL7 has revised and constantly rejected traditional data elements from specific local business needs, and only admits data elements that emerge from an analysis of what should generally be true about the entities of interest.

Why, then, does the RIM not seem to define the biomedical reality, e.g., molecular processes or disease processes, the ultimate subjects of the health information? To be sure, the RIM *does* provide structures (*form*) for modeling physical reality, but it leaves the definition of the *content* to specialty groups, either inside or outside HL7, and often in the form of terminology or ontologies. For example the HL7 drug knowledge model that is implemented with the U.S. FDA and pharmaceutical industry [8] contains classes for medicines, substances (Entities) and ingredient relationships (Roles relating the Entities). This is a general structure which supports the aggregation of a detailed and authoritative ontology of medicinal products. However, to describe the analysis of an unknown substance, or to determine the blood-level of a substance, HL7 uses laboratory measurement observation Acts instead of ingredient Roles, even though both are about the same chemical phenomenon, i.e., concentration. The Entity model is used for information that is known *a priori*, while observation acts are used for information that is only in the process of being discovered. This is why the clinical genomics model *seems* to conceptualize biomolecules as Acts, because it supports the process of identifying and discovering these molecules in particular instances.

HL7 also uses observation acts to describe that which may well be known *a priori*, but which is *not universally agreed* as to whether it is relevant or how it should be described. For example, we model the concentration of ingredients directly in the Roles but the color of medicines we model as observations. What seems like an arbitrary split

between information that is supposedly known *a priori* vs. discovered or whose conceptualization is generally agreed vs. unknown, is in fact the practical thing to do. The alternative would be to leave even simple generally agreeable *a priori* information in a generic undefined structure, as it is common in “Entity-Attribute-Value” models, which are omnipotent but have no normative power, or, to create a variety of competing and evolving detail models which would require frequent modifications, such as addition or removal of special attributes. The RIM has many structural similarities with existing EHR systems’ data models, featuring generic Observations structures, dynamically defined by data dictionaries, next to detailed *a priori* defined structures, which, when changed, would force expensive software updates. The modeling tradeoffs in the RIM are a reflection of the tacit knowledge and compromises of a community of practical system developers and users in health care.

A complete and integrated ontology of everything would certainly be nice to have; however, we think it is impractical and in fact dangerous to force such a model into being independently of the RIM. The moment such a model gained traction people would then expect that the RIM reflect that other model. Why should there be two models, if in the end one is to reflect the other? Instead, a single model of real world objects should suffice, but must contain well-defined features for information-management functions. These information-management functions are necessary to address the core problem of how intelligent agents (humans and heterogeneous software systems) arbitrate their perceptions of reality, how they communicate their intentions, and how they eventually effect changes to the real world. These are the core issue which HL7 addresses, and the remainder of this paper is dedicated to them.

3.3. Triple Standard – Act, Speech Act, and Documentation

We have introduced into the HL7 model the notion of *speech acts* [5], to describe the role that healthcare information plays in enabling cooperation. [6, 1] Speech acts are a generally accepted linguistic tool for understanding pragmatics, i.e., how language is used for achieving certain goals. Speech acts consist of *propositional content* and *illocutionary force* [7]. Propositional content is what the speech act *is about*; the illocutionary force is what it *accomplishes*. For instance, making assertions, demands, promises, proclamations, etc. accomplishes certain extra-linguistic effects in the behavior of others through illocutionary force. *Modality* is what the force accomplishes (e.g., assertion, demand, wish) subject to preparatory conditions, such as the relationship between the speaker and the one who is spoken to. Only if this relationship is appropriate the speech act can have its effect. For example, an order issued by an unauthorized person is invalid and will not have the intended effect. In the HL7 RIM, most of the features (attributes and participations) of the Act class (called “descriptive” features) carry the *propositional content*, and some features (called “inert” features) carry the information that substantiates the *illocutionary force*, such as the *mood-code* for modality, *author-participation* to substantiate the identity and role of the speaker and its relation to the receiver. Vizenor [2] took our speech-act analogy to further analyze cooperative healthcare actions, and he raises three major points of critique.

Firstly, Vizenor says that not all Acts in the RIM are speech-acts, e.g., an order for an injection is a speech act, but the injection itself is not a speech act, and that therefore the real act ontology and the speech-act ontology should be separately refined. We agree that the act of injecting is not a speech act, yet the medical record does not contain injections *per se*, but rather someone’s talking about injection as an order or a

report. The shared care record contains speech acts (even a simple assertive statement is a speech act) of which the physical act of injecting is simply reflected as propositional content. Trying to separate propositional content from its speech act is futile, because in the end, we have to represent both in a linguistic form. Propositional content is tied to speech just as one side of a coin is tied to the other.

Secondly, Vizenor says that the RIM documentations' considering Acts as "attributed statements" contradicts Acts being "speech acts", and that it is therefore wrong to tie attribution to the Act class. "Attribution" is the widely accepted practice to keep all data associated with its originator, and does not neutralize the illocutionary force of the speech acts. Instead, it simply substantiates the preparatory conditions that establish the illocutionary force in the first place. In human speech acts, the identity of the speaker is immediately given in the speech performance, the RIM only accounts for this. So, both propositional content and attribution are inseparable from speech acts.

Lastly, Vizenor identifies certain post-conditions in his analysis of speech acts, such as a promise creates an obligation of the *promiser* and a "claim" on the side of the *promisee*. Vizenor feels that records or documents represent these post-conditions, which he identifies as the "continuants" in his ontology and therefore expects to find them modeled as RIM Entities. The HL7 RIM, however, regards the post-conditions of speech acts as states inside the communicating systems that are sufficiently substantiated by the record of the speech act itself, which each system interprets given its aspect with regard to the speech act. This is why a "document" in the RIM is only a special form of speech act, in the form of a human readable text, while the functions of *documentation* and *recording* are supported by any RIM Act object.

4. Conclusion

Many formal technical criticisms as well as critique about the conceptual design of the RIM are rooted in misunderstandings, and more importantly in what seems to be a fundamental disagreement as to how a reference model for semantic and pragmatic interoperability should be created. We welcome the discussion on and off the record of published literature, as it helps to clarify the fundamental assumptions and opens up the RIM to much needed logical validation and improvement. However, we believe that if the fundamental approaches are at odds, the specific criticism will remain ineffective.

References

- [1] Schadow G, Russler DC, McDonald CJ. Conceptual alignment of electronic health record data with guideline and workflow knowledge. *Int J Med Inform.* 2001 Dec;64(2-3):259-74.
- [2] Vizenor L. Actions in health care organizations: an ontological analysis. *Medinfo.* 2004;11:1403-7.
- [3] Smith B, Ceusters W. HL7 RIM: an incoherent standard. *Medical Informatics Europe.* 2006. Submitted.
- [4] Ceusters W, Smith B. Strategies for referent tracking in electronic health records. *Biomed Inform.* 2005 Sep 27.
- [5] Austin JL. *How to do things with words.* Oxford: Oxford University Press, 1962.
- [6] Schadow G. *Krankenhauskommunikation mit HL7; Analyse, Implementation und Anwendung eines Protokollstandards für medizinische Datenkommunikation.* Aachen, Germany: Shaker-Verlag; 2000.
- [7] Loos EE, Anderson S, Day DH, Jordan PC, Wingate JD (Eds). *Glossary of linguistic terms.* <http://www.sil.org/linguistics/GlossaryOfLinguisticTerms> (Accessed on: 1/30/2004)
- [8] Schadow G, Gitterman S, Boyer S, Dolin RH eds. *HL7 v3.0 structured product labeling, release 2 [standard].* Ann Arbor, MI, Health Level Seven, 2005