From: *Studies in Health Technology and Informatics*, 2006; 124: 133–138
Presented at Medical Informatics Europe, Maastricht, August 2006.

HL7 RIM: An Incoherent Standard

Barry Smith[a,b,c,1] and Werner Ceusters[c]
[a] *Department of Philosophy, University at Buffalo, Buffalo NY, USA*
[b] *Institute for Formal Ontology and Medical Information Science, Saarbrücken, Germany*
[c] *Center of Excellence in Bioinformatics and Life Sciences and National Center for Biomedical Ontology, University at Buffalo, Buffalo NY, USA*

Abstract. The Health Level 7 Reference Information Model (HL7 RIM) is lauded by its authors as 'the foundation of healthcare interoperability'. Yet even after some 10 years of development work, the RIM is still subject to a variety of logical and ontological flaws, which has placed severe obstacles in the way of those who are called upon to develop implementations. We offer evidence that these obstacles are insurmountable and that the time has come to abandon an unworkable paradigm.
Keywords: HL7, RIM, standardization, ontology, realism.

1. Preamble

> A message to mapmakers: highways are not painted red, rivers don't have county lines running down the middle, and you can't see contour lines on a mountain. [1]

What follows is an exegesis and critique of the HL7 Reference Information Model (RIM). We will focus primarily on the relevant portions of [2], official codex of HL7 Version 3 ("Normative Edition 2005"), in which the RIM is asserted to be 'credible, clear, comprehensive, concise, and consistent', 'universally applicable', and 'extremely stable'. These assertions not only contradict many statements to be found within HL7's own internal email forums, they are also belied by the frequent revisions to which the RIM has been and is still being subjected, and by the fact that the RIM continues to be marked by major flaws (for which further documentation is provided at [3]). Such flaws include:

- *problems of implementation*: The decision by HL7 to adopt the new RIM-based methodology was adopted already in 1996; after ten years of effort, and considerable investment in the RIM itself and in the development of associated message types, DMIMs and RMIMs, the promised benefits of interoperability remain elusive;
- *problems of usability in specialist domains*: The RIM methodology consists in defining a set of 'normative' classes (Act, Role, and so on), with which are associated a rich stock

[1] Corresponding author: Barry Smith, Department of Philosophy, University at Buffalo, Buffalo NY 14260l, USA. Email: phismith@buffalo.edu; internet: http://ontology.buffalo.edu/smith.

of attributes. When the RIM is applied to a new domain (for example pharmacy), one needs to select from just these pre-defined attributes, rather as if one were attempting to create manufacturing software by drawing from a store containing pre-established parts for making every conceivable manufacturable thing, from lawnmowers to hunting bows. We do not know of even one example where a methodology of this sort has been made to work successfully;

- *problems of scope*: What are called 'Acts' (roughly: intentional actions) play an overwhelmingly central role in the RIM, so that of its 'foundation classes' only Entity (defined as: 'A physical thing, group of physical things or an organization capable of participating in Acts, while in a role') is left over to comprehend those things which are not Acts. How, on this basis, can the RIM deal transparently with information about, say, disease processes, drug interactions, wounds, accidents, bodily organs, and many other phenomena central to healthcare, given that the latter are neither Entities nor Acts?
- *problems of documentation*: The RIM documentation is not only disastrously unclear, and poorly integrated with those other parts of the V3 documentation for which the RIM itself is designed to serve as backbone; it is also subject to a series of internal inconsistencies (for example in its sloppy use of terms such as 'act', 'Act', 'Acts', 'action', 'ActClass' 'Act-instance', 'Act-object') of a sort which should surely be avoided in the context of work on messaging standards.
- *problems of learnability*: Can HL7 V3 be taught, and therefore engaged with and used by a wider public, given the amateurish quality of its documentation, the massive number of special cases, the frequent amendments, and the complex hurdles that must be overcome in creating a message?
- *problems of marketing*: Are the grandiose marketing claims made on behalf of HL7 V3 as 'the data standard for biomedical informatics' justifiable, given the many still unresolved problems on the technical side?

2. Double Standards

In spite of large investments in HL7-based information systems, some of them by national governments, and in spite of the now familiar difficulties encountered in creating working implementations on the basis of such investments, there is astonishingly little secondary literature on the HL7 standard itself. One consequence of the absence of evidence-based criticism from independent outsiders is the appearance of certain symptoms of intellectual inbreeding in HL7's own literature, which is marked above all by an air of boosting self-congratulation. We believe, against this background, that the HL7 endeavor can only benefit from forthright criticism, and it is in this spirit that our remarks are offered here.

For reasons of space, we shall focus on just one set of defects, which concerns the RIM's unsure treatment of the distinction between *information about an action* on the one hand and *this action itself* on the other. The former is what is *recorded* in a message or record. The latter is what *occurs*, for example within a hospital ward or laboratory. When challenged, RIM enthusiasts will insist that the RIM is concerned exclusively with the former – with information – and of course the very title of the RIM is in keeping with this

conception. Interspersed throughout its documentation, however, we find also many references to the latter, often conveyed by means of the very same expressions. This problem is of crucial importance, given the characteristic claim advanced on behalf of the RIM that it will provide a shared, rigorous semantics for HL7 V3 messages, of a type which V2 lacked. For it shows that what the RIM in fact provides is deep confusion.

As example, consider the treatment of the RIM class LivingSubject, which is defined, confusingly, as follows:

> A subtype of Entity representing an organism or complex animal, alive or not. (3.2.5)

Examples of this class are then stated to include: 'A person, dog, microorganism or a plant of any taxonomic group.' From this we infer that a person, such as you or me, is an example of a LivingSubject. At the same time in 3.2.1.13 (which is, oddly, the only subsection of 3.2.1 in [2]) we are told that LivingSubjects, including persons, can occupy just two 'states': *normal*, defined simply as 'the "typical" state', or *nullified*, defined as: 'The state representing the termination of an Entity instance that was created in error.'

Can it really be true that we are here being invited to postulate two kinds of death for human beings: normal death, and a special kind of death-through-nullification in the case of those persons who were created in error? Or is it not much rather the case that by 'LivingSubject' the RIM means not (as is asserted at 3.2.5) 'mammals, birds, fishes, bacteria, parasites, fungi and viruses' but rather *information about* such entities? The answer, bizarrely, is that *it means both of these things*; for there are in fact, co-existing side by side within its documentation, two distinct conceptions of what terms in the RIM are supposed to designate.

What we shall call the **information model conception of the RIM** is enunciated for example in 1.1 of [2]:

> The Health Level Seven (HL7) Reference Information Model (RIM) is a static model of health and health care information as viewed within the scope of HL7 standards development activities. It is the combined consensus view of information from the perspective of the HL7 working group and the HL7 international affiliates. The RIM is the ultimate source from which all HL7 version 3.0 protocol specification standards draw their information-related content.

The RIM, according to this first conception, is intended to provide a framework for the representation of the structures of and relationships between information that is independent of any particular technology or implementation environment. It is thus designed to support the work of database schema designers, software engineers and others by creating a single environment for messaging which can be shared by all healthcare institutions.

What we shall call the **reference ontology conception of the RIM** can be inferred, first, from the many programmatic statements describing the RIM's purpose, which is to facilitate consistent sharing and usage of data across multiple local contexts. For in striving to achieve this end of consistency (and thus to rectify problems affecting implementations of HL7 V2), the RIM cannot focus merely on healthcare messages themselves, as bodies of data. Rather it must provide a common benchmark for how such bodies of data are to be formulated by their senders and interpreted by their recipients. We can conceive of only one candidate benchmark for this purpose, namely the things and processes themselves within the domain of healthcare which messages are about and which are familiar to those engaged in message formulation. That this is indeed the benchmark adopted by the RIM will become clear when we examine the many passages in its documentation in which definitions and

examples are provided to elucidate the meanings of its terms.

Rather than distinguishing the two tasks, of *information model* and *reference ontology*, and addressing them in separation, the RIM seeks to tackle both simultaneously, through ambiguous use of language. Expressions drawn from the vocabulary of healthcare are used, alternately, both with familiar meanings (which enable the RIM to secure a necessary relation to corresponding activities in the healthcare domain) and with incompatible technical meanings (when the RIM is attempting to specify the relevant associated type of data). Thus for example 'stopping a medication' means both: *stopping a medication* and: *change of state in the record of a Substance Administration Act from Active to Aborted* (3.1.5).

3. 'Objects'

According to the information model conception, the RIM, and the HL7 messages defined in its terms, are about objects in information systems – hereafter called 'Objects' – which 'represent' things and processes in reality. Thus the human being named 'John Smith' is represented by an Object containing John Smith's demographic or medical data. This Object is different from John Smith himself. HL7's Glossary defines an 'Object' as:

> An instance of a class. A part of an information system containing a collection of related data (in the form of attributes) and procedures (methods) for operating on that data.

It defines an 'Instance' as: 'A case or an occurrence. For example, an instance of a class is an object.' Yet under the entry for the class 'Person' in the same Glossary we are told that 'Instances of Person include: John Smith, RN, Mary Jones, MD, etc.' For HL7, accordingly, 'John Smith, RN' is not the name of a human being; rather, it is the name of an Object which goes proxy for a human being in an information system.

Because Objects are bodies of data, and because different data about John Smith will be contained in the different information systems involved in messaging, 'John Smith, RN' will refer ambiguously to many John Smith Objects. How, then, is messaging about John Smith (the human being) possible, given that senders and recipients will associate distinct John Smith Objects with each given message?

This is, to be sure, a hard problem, and the HL7 community deserves some credit for having recognized its importance. The RIM can go part way towards alleviating it in the case of persons by insisting that information objects include corresponding unique identifiers such as social security numbers. Indeed the HL7 Glossary *defines* a 'Person' as a 'single human being who … must also be uniquely identifiable through one or more legal documents (e.g. Driver's License, Birth Certificate, etc.).' In the ideal case, at least (which would require each institution to maintain a mapping of its own locally unique patient IDs to all the locally unique patient IDs used by the institutions it is communicating with), such identifiers could serve to bind the different Objects together in such a way that they would all become properly associated with what we would normally think of as one and the same person. But what works for Persons will in almost every case not work for Objects representing things and processes of other types (for instance tumors, epidemics, adverse reactions to drugs), since unique identifiers for the latter are almost never available under present regimes for recording healthcare data [4].

4. 'Acts'

The term 'Act' refers, within those passages in the RIM which conform to the information model interpretation, not to acts (intentional actions), but rather to associated Objects. Thus at 3.1.1 'Act' is defined as meaning: 'A record of something that is being done, has been done, can be done, or is intended or requested to be done.'

Confusingly, however, we are provided, in explication of this definition, with examples not of *records* but of *intentional actions themselves* (referred to not as 'Acts' but as 'acts'):

> The kinds of acts that are common in health care are (1) a clinical observation, (2) an assessment of health condition (such as problems and diagnoses), (3) healthcare goals, (4) treatment services (such as medication, surgery, physical and psychological therapy), (5) assisting, monitoring or attending, (6) training and education services to patients and their next of kin, (7) and notary services (such as advanced directives or living will), (8) editing and maintaining documents, and many others.

At 3.1.14, similarly, the class PatientEncounter (a subclass of Act) is defined as 'An interaction between a patient and care provider(s) for the purpose of providing healthcare-related service(s).' An instance of PatientEncounter, then, is not a *record* of an action, but this action itself – as is made clear by the examples provided to elucidate this definition, which include: emergency room visit, field visit, occupational therapy. The class Procedure is defined as 'An Act whose immediate and primary outcome (post-condition) is the alteration of the physical condition of the subject' (3.1.15), with examples: chiropractic treatment, balneotherapy, acupuncture, straightening rivers, draining swamps.

The class Observation is defined as:

> An Act of recognizing and noting information about the subject, and whose immediate and primary outcome (post-condition) is new data about a subject. Observations often involve measurement or other elaborate methods of investigation, but may also be simply assertive statements. (3.1.13)

From this we can infer, dizzyingly, that we are again dealing not with *records* of actions (of observing, measuring, etc.), but rather with these actions themselves, a view supported also by the RIM's assertion, still in 3.1.13, to the effect that Observations 'are professional acts … and as such are intentional actions', as also by its treatment of other subclasses of Act, such as DeviceTask, SubstanceAdministration, and Supply – all of which are similarly interpretable only against the background of a reference ontology conception of the RIM and thus as standing in conflict with the RIM's own definition of 'Act'. When, in contrast, we turn to other subclasses of Act, for example Account, or FinancialTransaction, then we find that the corresponding definitions revert to the information model conception.

We are told at 3.1.1 that, while an Act-instance 'represents a "statement"', 'the Act class is this attributable statement'. We can interpret this strange language as follows: each Act-instance is a statement describing what some clinician on some occasion has heard, seen, thought, or done. The Act class is the class of such coded, attributable statements. An Act-instance is, more precisely (in 3.1.1 at least), either an attributable statement or some similar attributable use of language such as an order or request, which serves as the coded accompaniment – the record – of what takes place in reality, for example when a surgical procedure is performed. There then follows a most peculiar passage, which has been criticized already in [5]:

> Act as statements or speech-acts are the only representation of real world facts or processes in the HL7 RIM. The truth about the real world is constructed through a combination (and arbitration) of such attributed

> statements only, and there is no class in the RIM whose objects represent "objective state of affairs" or "real processes" independent from attributed statements. As such, *there is no distinction between an activity and its documentation*. (3.1.1, emphasis added)

The italicized passage captures what we might think of as the naked essence of the confusion at the heart of the RIM between reference ontology and information model.

5. Conclusion: Can the RIM be Saved?

Messaging standards are like telephone systems: they can function sensibly only if there is a large network of willing users. This means that to succeed such artifacts must be marked by *clear use of language* and *clear documentation*. The RIM documentation, as we have seen, is systematically ambiguous. What is needed, if HL7 V3 is to satisfy its need for a uniform information representation that is coherent, clear and implementable, are two separate, though of course related, artifacts, which might be called 'Reference Ontology of the Healthcare Domain' and 'Model of Healthcare Information', respectively. The former would include those categories, such as *thing*, *process*, *anatomical structure*, *disease*, *infection*, *procedure*, etc., needed to provide a compact and coherent high-level framework in terms of which the lower-level types captured in vocabularies like SNOMED CT could be coherently organized [6]. The latter would include those categories, such as *message*, *document*, *record*, *observation*, etc., needed to specify how information about the entities that instantiate the mentioned types can be combined into meaningful units and used for further processing. HL7's Clinical Document Architecture could then be related in an appropriate way to this Model of Healthcare Information, thereby avoiding the current counterintuitive stopgap, which forces a document to be an Act. And HL7's Clinical Genomics Standard Specifications could similarly be related in an appropriate way to the Reference Ontology of the Healthcare Domain, thereby avoiding the no less counterintuitive current stopgap, which identifies an individual allele as a special kind of Observation.

Acknowledgements: This paper was written under the auspices of the Wolfgang Paul Program of the Humboldt Foundation, the EU Network of Excellence SemanticMining, the Volkswagen Foundation, and the National Center for Biomedical Ontology (the latter funded by the National Institutes of Health through the NIH Roadmap for Medical Research). Thanks for helpful comments are due also to Thomas Beale, Walter Dierckx, Amnon Eden and Gunther Schadow, who however bear no responsibility for the positions here adopted.

References

[1] Kent W. Data and Reality, Bloomington IN, 2000, p. v.
[2] ANSI/HL7 V3 RIM, R1-2003: RIM Version V 02-11. Membership Normative Ballot Last Published: 11/22/2005 8:05 PM.
[3] http://ontology.buffalo.edu/HL7.
[4] Ceusters W, Smith B. Strategies for referent tracking in Electronic Health Records, J Biomed Inform, in press.
[5] Vizenor L. Actions in health care organizations: An ontological analysis. Proc MedInfo 2004, San Francisco, 1403-1407.
[6] Rosse C, Kumar A, Mejino JLV, Cook DL, Detwiler LT, Smith B. A strategy for improving and integrating biomedical ontologies. Proc AMIA Symp 2005;:639-643.