

Human Action in the Healthcare Domain: A Critical Analysis of HL7's Reference Information Model*

Barry Smith

Lowell Vizenor

Werner Ceusters

Abstract

If we are to develop efficient, reliable and secure means for sharing information across healthcare systems and organizations, then a careful analysis of human actions will be needed. To address this need, the HL7 organization has proposed its Reference Information Model (RIM), which is designed to provide a comprehensive representation of the entire domain of healthcare centered around the phenomenon of human action. Taking the Basic Formal Ontology as our starting point, we examine the RIM from an ontological point of view, describing how it fails to provide a representation of the healthcare domain which would enjoy the sort of clarity, coherence, rigor and completeness that is claimed on its behalf.

1. Introduction

Information and communication technology has not only altered the way that medical information is generated, stored, analyzed, and shared across and within healthcare organizations, it has also come to be associated with the promise that it will increase the safety, efficiency and cost-effectiveness of healthcare. The hope is that the electronic health record (EHR) and associated reporting, analysis and decision support technologies will facilitate the diffusion and dissemination of healthcare information, thereby allowing the systematic use of clinical guidelines

* We dedicate this paper to Ingvar Johansson our friend and colleague at the Institute for Formal Ontology and Medical Information Science (IFOMIS) from 2002 to 2008. The interdisciplinary research group at IFOMIS included representatives from ontology, biomedical informatics, and linguistics, and Ingvar himself contributed in important ways in all of these areas — including the peculiar interaction between ontology, medicine and speech act theory which forms our topic in what follows.

and outcomes measures in ways that will bring benefits to human health in the form of increased safety, effectiveness and economy.

It has for long been clear, however, that many difficulties must be overcome before the promise of health informatics can be fully harvested. One such difficulty arises from the fact that the single doctor — single patient nexus has been largely superseded by a regime in which the typical patient is managed by a team of health care professionals, each specializing in one aspect of care. This is significant because different departments within a healthcare organization have different disciplinary cultures and employ different terminologies and data formats to talk about what are putatively the same phenomena. Current efforts to develop efficient means for sharing information across healthcare systems and organizations that have some prospect of overcoming this and a range of similar problems must find effective ways to share information in an intuitive and stable way that ensures that meaning is preserved.

2. The Birth of the HL7 RIM

To see how difficult a task this is, we describe one ambitious and highly influential effort to standardize healthcare information across the entire domain of healthcare that has been advanced by the Health Level 7 (HL7) organization, one of several ANSI-accredited Standards Developing Organizations operating in the healthcare arena.¹

Already in the 1990s HL7 enjoyed considerable success through its creation of a widely used standard for healthcare messaging which was established as mandatory for communication between US Federal Government-funded healthcare organizations. This standard is now commonly referred to as HL7 version 2 — or v2, for short.² v2 enabled healthcare applications to exchange clinical, demographic and administrative data in digital form on the basis of what we can think of as a walkie-talkie paradigm. Significantly, the v2 standard was designed to meet the interface requirements of the healthcare system in its entirety rather than focusing on the requirements of just one area of healthcare such as pharmacy, imaging services or insurance claims management.

¹ <http://hl7.org>, last accessed November 12, 2012.

² We here ignore the differences between successive sub-versions of the v2 standard.

Unfortunately, as v2 proved ever more popular, it also led to the creation of manifold v2 dialects, which over time brought about a situation in which messaging interoperability was maintained in many cases only within, and not between, healthcare organizations. Fatefully, the response of the HL7 organization to this problem was to develop, starting in the 1990s, an abstract model of the entire domain of patient care called the Reference Information Model (RIM) that was intended to serve as a unified framework for the sharing of information and the usage of data across the entire domain of healthcare and to serve as a constraint on all subsequent HL7 standards. By regulating in this way what would be allowed to be communicated via subsequent v2-style HL7 messaging systems, v2's problems of dialect formation would, it was hoped, be solved.

3. HL7's Act-Centered View of Healthcare

The RIM starts out from the assumption that any profession or business, including healthcare, can be viewed as consisting primarily of a series of *intentional actions* on the part of responsible actors working within an organizational framework. The varieties of such intentional actions relevant for healthcare include:

- assessment of health conditions (such as the taking of your pulse, or the weighing of your baby);
- provision of treatment services (such as performing surgery, or administering drugs);
- informing of patients and their next of kin about health conditions;
- provision of notary services (such as the preparation of a living will);
- editing and maintaining of documents;
- ordering and accepting delivery of supplies;
- reporting to government agencies;
- billing;
- and many more.

Interestingly, HL7 explicitly acknowledges the influence of philosophy in its creation of the RIM, whose act-centered view of healthcare draws for its underlying framework on the speech act theory developed by J. L. Austin in Austin (1962) and by John R. Searle in Searle (1969).

Austin and Searle were among the first philosophers to recognize the theoretical significance of the fact that *what we can do with words* goes well beyond uses of language of the statement-making sort.¹ We can make requests, issue commands, make promises, ask questions, and so on, and actions of these sorts are marked by the fact that the very utterance of words brings about some extra-linguistic result, as for example when the making of a promise brings about the result that the maker of the promise stands under a certain obligation to perform a certain act.

Speech acts of different types, now, can share the same propositional content. Thus I can command that you open the door; I can suggest that you open the door; I can ask whether you will open the door; and so on. Moreover, as was recognized by pioneers of the logic of action,² there is a sense in which this same propositional content can be shared also by actions of a non-linguistic sort, as when a command is obeyed or an instruction is followed. In virtue of this sharing of contents speech acts and other human actions form certain standard sorts of sequences, as when, for example, a question is followed by an answer, an act of transmitting information is followed by an act of acknowledgement, a promise is followed by the performance of the promised act, or an act of ordering bedding supplies is followed by acts of, for example, transporting, receiving, billing and paying for the bedding supplies delivered.

The RIM sees this idea of shared propositional content as an attractive way of modeling how the domain of human actions is organized. In sequences such as the sort described, we are to utilize the same RIM ‘classCode’ — for example ‘replenish bedding supplies!’ or ‘register this patient!’ or ‘administer this drug!’ — to capture the common content of what is involved in each successive act within the sequence, combined with a succession of different ‘moodCodes’ (such as ‘order’ or ‘command’) to capture what is peculiar about each succeeding act. The ‘model of healthcare information’ that is created in this way is seen by the authors of the RIM as providing an efficient and reliable framework for ensuring successful communication of meaning within and between healthcare information systems.

Of course there are many features of healthcare that go beyond the category of action. These include the *participants* of the actions them-

¹ For a broader view of the history of speech act theory see Barry Smith (1990).

² See for example G. H. von Wright (1963).

selves, both agent and patient. They include the *roles* these participants play in actions, including their *authority* to perform given actions. They include the sorts of entities to which these actions give rise (Smith 2003), such as *obligations*, *claims* and *electronic documents*. They include *diseases*, and the associated causal processes inside the organism, including processes such as birth and death. They include *material objects* such as pharmaceutical products, DNA samples, equipment and buildings, and they include *organizations and institutional entities* such as insurance companies, government agencies, and laboratories. All of these must be taken into consideration in a complete ontology of the healthcare domain. Ultimately, an ontologically adequate language for communication of healthcare information should have the resources to capture all of the items on this list and all of the different sorts of relations that hold between them, and to do this in a maximally intuitive way that is at the same time easily modifiable as the needs and practices of healthcare organizations change with time.

4. The RIM Straightjacket

HL7's current documentation of the RIM standard appears in an International Standards Organization (ISO) document entitled "Health informatics: — HL7 version 3 — Reference information model — Release 4", and described as a "Draft International Standard (DIS)".¹ The document serves as the basis for a ballot to establish the new Release 4 of the RIM as an ISO standard.²

As the document makes clear, the RIM requires that all healthcare information will be organized in terms of just the six "backbone" classes presented in Table 1 below. As the reader will see, this Table comprises two lists, of descriptions, and of definitions, which seem (to us, at least) to be in various ways inconsistent, even though they are taken from a single document. It is difficult to write clearly about HL7 specifications when these specifications are themselves formulated in inconsistent ways. Moreover, there are a number of further problematic issues with the individual entries in these two lists.

¹ Health informatics: – HL7 version 3 – Reference information model – Release 4 (Document: ISO/HL7 21731:2011(E)), <http://standardsproposals.bsigroup.com/Home/getPDF/1361>, last accessed November 28, 2012.

² As we understand the matter, the proposal is to establish the RIM as an international standard for being the RIM.

First, they are marked by an embarrassing circularity, as for example in the definitions of ‘Entity’ and ‘Role’, which can be of no possible assistance to someone who does not already understand what HL7 takes to be the meaning of the terms defined.

Second, are problems of ambiguity, for example when we are told that *Act* ‘represents the actions that are executed and must be documented as health care is managed and provided’ does this mean that actions voluntarily recorded do not fall under the heading of Act? Clarity, here, would demand a distinction, to which we shall return below, between health-care actions in general and actions of documentation in particular.¹

| Backbone Class | Description (from 0.2: “RIM as an abstract model”) | Definition (from 9.3: “Code System”) |
|------------------------|---|--|
| <i>Act</i> | represents the actions that are executed and must be documented as health care is managed and provided | a record of something that is being done, has been done, can be done, or is intended or requested to be done. |
| <i>Participation</i> | expresses the context for an act in terms such as who performed it, for whom it was done, where it was done, etc. | indicates that the target of the participation is involved in some manner in the act, but does not qualify how. |
| <i>Entity</i> | represents the physical things and beings that are of interest to, and take part in health care | a physical thing, group of physical things or an organization capable of participating in Acts while in a role. |
| <i>Role</i> | establishes the roles that entities play as they participate in health care acts | a competency of the Entity that plays the Role as identified, defined, guaranteed, or acknowledged by the Entity that scopes the Role. |
| <i>ActRelationship</i> | represents the binding of one act to another, such as the relationship between an order for an observation and the observation event as it occurs | a directed association between a source Act and a target Act. |

¹ Thus although HL7 is uncertain in its habit in this respect, we shall for purposes of clarity normally capitalize the first letter of HL7 terms such as ‘Act’, ‘ActClass’, and so forth. When terms such as ‘entity’, ‘act’, ‘action’, and ‘role’ appear without initial capitals in this essay (other than in quotations from HL7 documents), then their common meanings are intended.

| | | |
|-----------------|---|--|
| <i>RoleLink</i> | represents relationships between individual roles | a connection between two roles expressing a dependency between those roles and permitting the authorization or nullification of a dependent role based on status changes in its causal or directing role. The RoleLink may be operated over time and thus whose state and identity must be managed [<i>sic</i>]. |
|-----------------|---|--|

Table 1: The Six Backbone Classes of the HL7 Reference Information Model

Unfortunately the RIM's definition of Act positively undermines a distinction along these lines by identifying 'Act' with 'Record' — and this, even though the definitions of Entity and Role provided in the same ISO document see the latter not as records, but rather as the Entities and Roles themselves, thus further consolidating the 'incoherence' which we identified in the RIM already in 2006 (Smith and Ceusters 2006).¹ Third, are problems of interpretability, as when the document oscillates — to us mysteriously — between the use of 'act' and 'Act', or 'role' and 'Role', sometimes within one and the same sentence.²

Fourth, and perhaps most importantly, is the apparent narrowness of scope (in a standard 'reference information model' that is intended to cover the entire domain of healthcare). The RIM's list of backbone classes is intended to be exhaustive, yet important families of items seem to be excluded. Above all, where is the place for diseases and for disease processes inside the patient's body? Where is the place for hospital-related adverse events such as falls or spills or leakages of radioactive materials? These are not Acts, they are not contexts for Acts, they are not Entities, and they are not Roles. Where, then, do they fit within the RIM?³

In what follows we discuss our attempt to make sense of the RIM's backbone classes in terms of what they *include*, and drawing on HL7's own documentation and usage. We exploit in this connection certain fundamental ontological categories distinguished by philosophers in

¹ See Barry Smith and Werner Ceusters (2006: 133–138).

² See for example the definition of 'Disciplinary action' as: 'An action taken with respect to a subject Entity by a regulatory or authoritative body with supervisory capacity over that entity. The action is taken in response to behavior by the subject Entity that body finds to be undesirable.'

³ We discuss HL7's response to this question in Smith and Ceusters (2006), *op. cit.*

dealing with speech acts and similar phenomena, categories which reflect, within a systematic, logical framework, the central commonsensical distinctions such as that between a thing and an event, or between what is particular and what is general. We believe that any framework of definitions that cannot be cashed out intelligibly in terms of such distinctions will not be teachable to, and learnable by, normal human beings, and thus will likely lead to errors and confusions (and thus to the very sorts of inconsistent development which were responsible for HL7's problems of dialect formation).¹

For this purpose we utilize as our instrument of evaluation the Basic Formal Ontology (BFO), an upper-level ontology originally developed in IFOMIS and now used by many groups of researchers throughout the world as a vehicle for promoting interoperability of systems designed for handling scientific and many other sorts of data.² We select BFO as framework for our assessment of the RIM because it contains a set of categorizations which have been both well-tested from many different perspectives and also carefully defined and elucidated from a logical point of view.

5. BFO: Independent Continuant

We use 'entity' (with lower case 'e'), in what follows, as an ontological term of art comprehending all items (objects, things, features, attributes, patterns,...) that exist in any way. (HL7's 'Entity' is thus much narrower in its extension.) All real-world entities, from the BFO perspective, for example all entities of the sort that we encounter in the domain of healthcare, fall into one of two exclusive categories of *continuant* and *occurrent* (Grenon, Smith and Goldberg 2005).

Continuants are entities which continue to exist through time; they preserve their identity from one moment to the next even while undergoing a variety of different sorts of changes. Continuants are divided by BFO into the two sub-categories of *independent* and *dependent continuant*, the latter being distinguished by the fact that they depend for their existence on the former in the way in which, for example, the tem-

¹ "Are the ISO 21090 Data Types Too Complex?", <http://hl7-watch.blogspot.com/2010/11/are-iso-21090-data-types-too-complex.html>, last accessed November 12, 2012. This is just one example of multiple posts at this site documenting the RIM's unteachability.

² <http://www.ifomis.org/bfo/>, last accessed November 12, 2012.

perature or mass of a material body dependence on this material body for its existence.

Typical examples of independent continuants from the healthcare domain include human beings, buildings, wheelchairs, scalpels, and paper documents in filing cabinets. Each of these entities continues to exist through time even as it undergoes changes, for example, a human being will continue to exist, and preserve its identity, even as it grows and ages over time.

The RIM's Entity seems, at first sight, to be a close analogue to what BFO identifies as independent continuants. Entities are described by the RIM as 'physical things and beings that are of interest to, and take part in health care'. However, when we examine some of the subtypes of Entity in the RIM, as illustrated in Table 2, we find a number of items which are not physical in the normal meaning of the word.

Most blatant is 'Imaging Modality', which (in conformity with standard usage among radiologists) is asserted to be a subclass of 'Device'.¹ It is, however, *defined* by the RIM as: 'Class to contain unique attributes of diagnostic imaging equipment'. This is to confuse a piece of equipment with one or more of the attributes of a piece of equipment. This confusion may well be compatible with the RIM's *description* of Entity (see Table 1), if the latter is to be read as having '(physical entities) and (beings)'² as its intended scope and if we are allowed to include attributes as 'beings'. But it is incompatible with the RIM's *definition* of 'Entity', since an attribute of a physical thing is neither a physical thing, nor a group of physical things, nor an organization capable of participating in Acts while in a role.

Worryingly, this example strongly suggests also that, when HL7's authors are formulating definitions, they have no sure understanding of the meaning of the very word 'definition'. In the case of 'Organization' we are indeed provided with a statement of necessary and sufficient conditions and thus with a definition of a logically recognizable sort. In the case of 'Imaging Modality', however, we are provided with something like a statement of the reasons why those responsible for introducing a term thought it necessary to do so (as if one were to define 'screw-

¹ <http://hl7.org>, last accessed November 12, 2012.

² As contrasted with 'physical (attributes and beings)'.

Entity =def. A physical thing, group of physical things or an organization capable of participating in Acts while in a role.

Living Subject =def. Anything that essentially has the property of life, independent of current state (a dead human corpse is still essentially a living subject).

Health Chart Entity =def. A health chart included to serve as a document receiving entity in the management of medical records.

Organization =def. A social or legal structure formed by human beings.

Group =def. A grouping of resources (personnel, material, or places) to be used for scheduling purposes. May be a pool of like-type resources, a team, or combination of personnel, material and places.

Place =def. A physical place or site with its containing structure. May be natural or man-made. The geographic position of a place may or may not be constant.

Material = def. Any thing that has extension in space and mass, may be of living or non-living origin.

ManufacturedMaterial =def. An Entity or combination of Entities transformed for a particular purpose by a manufacturing process.

Device =def. A subtype of ManufacturedMaterial used in an activity, without being substantially changed through that activity.

Certificate Representation =def. A physical artifact that stores information about the granting of authorization.

Imaging Modality =def. Class to contain unique attributes of diagnostic imaging equipment.

Table 2: HL7 RIM Entity class and selected subclasses¹

driver’, for example, as meaning ‘a term I will need next week when I reach the letter “S” in my list’).

Another problem example in Table 2 is ‘Health Chart Entity’, an immediate subclass of RIM ‘Entity’ that is defined as follows:

Health Chart Entity =def. A health chart included to serve as a document receiving entity in the management of medical records.

¹ Health informatics: – HL7 version 3 – Reference information model – Release 4 (Document: ISO/HL7 21731:2011(E)), *op. cit.*

Because the definition is circular, there is no easy way to understand its meaning. Is a Health Chart Entity a person or agency that can ‘receive’ a document? If so, then why is it not included as a child of ‘Person or Organization’. Or does ‘Health Chart Entity’ refer to the document itself? If so, then this would bring the implication that such a document is, in accordance with the RIM’s definition of ‘Entity’, as a ‘physical thing, group of physical things or an organization capable of participating in Acts while in a role’. To conceive a document as a physical thing, however, creates problems in view of the fact that the documents of interest to a healthcare organization will in many cases be electronic documents, and thus information artifacts (abstract patterns created through special processes which may be stored simultaneously on many different devices). As we have argued elsewhere, this is an item of a sort which calls for a treatment quite different from that of physical entities, a treatment that must distinguish clearly between the device that stores information and the information entities that are stored.¹ That the RIM does not acknowledge this distinction is seen in its treatment of ‘certificate representation’, which is defined as ‘A physical artifact that stores information about the granting of authorization’ and is asserted at the same time to be a subtype of ‘device’, defined as a subtype of material, and thus as ‘having extension in space and mass’. What is the *mass* of an e-certificate granting authorization?

Further problems are raised by the class ‘living subject’, which is for some reason *not* treated as a child of ‘entity class material’, even though the latter is defined by the RIM as ‘Any thing that has extension in space and mass, may be of living or non-living origin’. The problems here are compounded still further when we are told that ‘a dead human corpse is still essentially a living subject’.

6. BFO: Dependent Continuant and the RIM: RoleClass

Dependent continuants, in BFO, are the *states*, *properties*, *qualities*, and *roles* of patients, administrators and so forth. The category of dependent continuant is particularly important for an understanding of the ontology of social reality. Examples of special relevance for us here are the *mental* and *normative* states to which some actions give rise, including

¹ <http://code.google.com/p/information-artifact-ontology/>, last accessed November 12, 2012.

above all the *intentions* of the participants on the one hand and their *obligations* and *claims* on the other. The category of dependent continuants includes also the *capacities* and *skills* (counted as *dispositions* in BFO) of healthcare personnel, such as the ability to speak Spanish or to perform complex medical procedures; the *roles* — for example the nurse or patient roles — that participants play in actions; and their *authority* to perform given actions in virtue of having these roles. These entities are continuants in the sense that, like organisms and molecules, they preserve their identity over time. For example, an intention is a state; that is, it is something that endures from point of inception to point of realization. An entity of this sort is *dependent* in the sense that it requires the support of at least one other entity — its bearer — in order to exist. A *relation of authority* is similarly a dependent continuant, in this case of a sort that has a multiplicity of bearers, namely (i) the human being who has the authority in question, and (ii) the human being(s) over whom this authority is wielded.

Some dependent continuants are captured in the RIM by the classes *Role* and *RoleLink*. In the RIM, an Entity which participates in an Act must do so in a particular Role. The Role defines the Entity's competency (which actions it can perform) and constraints (which actions it cannot perform). In some cases, the Role connects the player of the Role to those bodies, groups, or agencies that have the power to recognize the Role. An example from the RIM is *LicensedEntityRole*, which is a relationship in which, for example, a medical authority certifies a medical caregiver as being permitted to perform certain activities that fall under the jurisdiction of the medical authority in question. The RoleLink class defines connections between Roles. Examples include *has direct authority over* and *has indirect authority over*. In this way the RIM comprehends chains of authority in an organization. While there are some similarities between HL7 Roles and BFO dependent continuants, however, there are also significant differences, to one of which we now turn.

From the BFO point of view, the universals represented in an ontology are instantiated by particulars. Universals correspond to the general terms (such as 'cell' or 'electron' or 'pneumonia') used in scientific texts and also to the even more general terms (such as 'independent continuant' and 'entity') used in high-level, domain-neutral ontologies such as BFO that have been designed to support the data annotation and

cross-domain data integration needs of scientists. Particulars are the entities that we can observe, for example in the lab or clinic, or record, for example when we register a baby's weight or a nurse's promotion.

Particulars then instantiate corresponding universals (Smith and Ceusters 2010). Just as humans, hospitals, kidneys, and so forth, instantiate the universals *person*, *hospital* and *kidney*, so dependent continuants such as Pippa's weight, or John's nurse role, or Mary's authority over Harvey, instantiate universals such as *nurse role* and *authority*, respectively. Just as it is common in a healthcare setting to find multiple instantiations of the universal *human*, so it is common to find multiple instantiations of the same *role*. For example, *nurse practitioner role* is multiply instantiated whenever a hospital has more than one nurse practitioner; but in each case it is the same role universal that is being instantiated. In the case of an authority role, the picture might look like this:

| | independent continuant | dependent continuant |
|------------|------------------------|------------------------------|
| universal | <i>organism</i> | <i>role</i> |
| | <i>person</i> | <i>authority role</i> |
| particular | Mary, Harvey | Mary's authority over Harvey |

Between Mary and the universal *person*, and between Mary's role as someone who has this specific authority in this specific healthcare institution, and the universal *authority* there obtains the relation of *instantiation*. Between Mary's authority and Mary herself there obtains the relation of *inherence* (Mary herself is the *bearer* of her authority). Between *person* and *organism* and between *nurse role* and *role* there obtains the relation of *subsumption* (*person* is a subtype, or sub-universal, of *organism*).

Can we now identify 'Role' in the RIM with what is called 'role' in BFO? Unfortunately not. Again, the extension of Role has been conceived by the authors of the RIM in a seemingly arbitrary way, with the result that it is too ontologically heterogeneous to identify with any class of entities that has been coherently defined. Certainly some subclasses of

HL7 Role represent dependent continuants in BFO's sense, as for example in cases such as this:

| | |
|---------------------|--|
| Healthcare provider | An entity (player) that is authorized to provide health care services by some authorizing agency (scoper). |
| Patient | A Role of LivingSubject (player) as a recipient of health care services from a healthcare provider (scoper). |
| Health cart | The role of material (player) that is the physical health chart belonging to an organization (scoper). |

Table 3: Examples of subclasses of HL7 Role¹

But there are on the other hand also specializations of Role which have nothing to do with roles as commonly conceived. Consider what the RIM calls *RoleClassOntological*, including *has generalization*, *instance* and *subsumed by* as subtypes (see Table 4).² The latter, unfortunately, have nothing to do with *roles* in the commonsensical usage of this term.

| | |
|----------------------|--|
| RoleClassOntological | A relationship in which the scoping Entity defines or specifies what the playing Entity is. Thus, the player's "being" (Greek: ontos) is specified. [Description: RoleClassOntological is an abstract domain that collects roles in which the playing entity is defined or specified by the scoping entity.] |
| Subsumed by | Relates a prevailing record of an Entity (scoper) with another record (player) that it subsumes. Examples: Show a correct new Person object (scoper) that subsumes one or more duplicate Person objects that had accidentally been created for the same physical person. Constraints: Both the player and scoper must have the same classCode. |
| Has generalization | Relates a specialized material concept (player) to its generalization (scoper). |
| Instance | An individual piece of material (player) instantiating a class of material (scoper). |

¹ Health informatics: – HL7 version 3 – Reference information model – Release 4 (Document: ISO/HL7 21731:2011(E)), *op. cit.*

² RoleClassOntological is introduced by the RIM as an immediate subclass of Role, but in a separate group from Patient, Caregiver, and other roles, alongside RoleClassAssociative and RoleClassPartitive.

| | |
|-------------------|---|
| Equivalent entity | Specifies the player Entity (the equivalent Entity) as an Entity that is considered to be equivalent to a reference Entity (scoper). The equivalence is in principle a symmetric relationship, however, it is expected that the scoper is a reference entity which serves as reference entity for multiple different equivalent entities. Examples: An innovator's medicine formulation is the reference for "generics", i.e., formulations manufactured differently but having been proven to be biologically equivalent to the reference medicine. Another example is a reference ingredient that serves as basis for quantity specifications (basis of strength, e.g., <i>metoprolol succinate</i> specified in terms of <i>metoprolol tartrate</i> .) |
| Same | The "same" role asserts an identity between playing and scoping entities, i.e., that they are in fact two records of the same entity instance, and, in the case of discrepancies (e.g. different DOB, gender), that one or both are in error. Usage: playing and scoping entities must have same classCode, but need not have identical attributes or values. Example: a provider registry maintains sets of conflicting demographic data for what is reported to be the same individual. |

Table 4: The RIM's RoleClassOntological and examples of its subtypes¹

The RIM's authors have, it would seem, for some reason come to the conclusion that it is necessary to add basic ontological terms and relations (such as 'kind', 'instance', 'subsumed by', and so forth) into the framework of the RIM. The recommended approach, in such circumstances, would be to adopt the best practices common in the appropriate discipline, which is in this case the discipline of ontology. (And, for all their differences, there is a great deal of commonality among the three most widely used upper-level ontologies — namely BFO, DOLCE and SUMO — as concerns basic ontological terms and relations.) Unfortunately, as so often been in the history of HL7, the RIM's authors chose to develop a special, idiosyncratic framework of their own, even though this meant sacrificing interoperability with externally developed standards. Moreover, they chose to do this on the basis of the assumption that the terms and relations in question must be subsumed under the RIM's existing six backbone classes.

¹ Health informatics: – HL7 version 3 – Reference information model – Release 4 (Document: ISO/HL7 21731:2011(E)), *op. cit.*

The result, we are sorry to say, is an egregious potage of confusion. To view *subsumed by* as a role is analogous to viewing the relation between, say, the species *rabbit* and the genus *mammal* by conceiving the latter as a role played by the former. And worse: all of the RIM's backbone classes, and indeed the entire structure of the RIM, require for their understanding the prior distinction between particular and universal. It does not make sense to conceive this distinction itself — under the heading 'Instance' — in terms of Roles.

7. BFO: Occurrent and the RIM ActClass

Occurrents (also called events or processes) are defined from the BFO point of view as being items which are such that they unfold themselves in their successive phases. Thus in contrast to continuants, occurrents never exist in full in any single instant of time.¹ The life of a patient is an occurrent, as also is the course or history of a given disease or of a given treatment. Actions are occurrents, and so also are sequences of actions, from planning, to issuing of orders, to the execution of a plan. (Plans themselves however are continuants: thus they endure continuously through time until they reach the point of complete execution or abandonment.)

Act represents the closest analogue in the RIM to occurrents, but it is at best a weak analogue, since as we have argued elsewhere (see Smith and Ceusters 2006), it is defined both too narrowly and too broadly, to a degree that gives rise to the hypothesis that (as in the case of 'Role') it is not capable of being coherently defined. Here we pursue the assumption that — in keeping with the conception of the RIM as a 'model of health-care *information*', rather than of the *reality* which such information is about — *Act* does not comprehend actions themselves, but rather only the records which arise when actions are documented by a healthcare professional in either a clinical or an administrative context. These records themselves are (for the RIM) Acts. (That such an identification brings confusion to the user is revealed most poignantly in the fact that the RIM / ISO document referred to above itself contains multiple passages, including the definition of Entity quoted in our Table 1, in

¹ We leave aside here the case of occurrent boundaries, for example beginnings and endings.

which ‘Act’ is used to mean, not ‘the record of an act’, but rather the act itself, in which some Entity participates.)

To help us to understand the rationale behind the RIM’s design, it is useful to point to the view of the medical record on which it is based, a view formulated by Rector and Nolan in 1991, according to which the medical record is a record, not of *what is the case on the side of the patient*, but rather of ‘what clinicians have said about what they have heard, seen, thought and done’ (Rector, Nolan and Kay 1991). From this point of view, not faithfulness to the clinical history and care of the patient is the fundamental criterion for what gets included in record, but rather, as Rector and Nolan express it:

The first consequence of our view of faithfulness is that the information in the medical record itself is not about what was “true” of the patient but what was observed and believed by clinicians.¹

Consider the case where physician A documents (at time t and place p) that physician B obtains a blood sample from patient C. Whether or not physician A actually documented that ‘physician B obtains a blood sample from patient C’ is of vital importance to the medical record. What is of lesser importance, according to the Rector-Nolan view, is whether or not the proposition ‘physician B obtains a blood sample from patient C’ is true. The idea is that information about the real world can be brought into the medical record only through records — *descriptions of acts* — formulated by suitably authorized persons. Since these descriptions are always attributed to someone, it is possible to have a medical record which contains statements about one and the same real world activity which yet disagree in their propositional content (for example because their authors were provided with conflicting information). The medical record will then still be consistent, for it will contain descriptions not of the form

physician B obtained a blood sample from patient C

physician B did not obtain a blood sample from patient C

but rather of the form

¹ On the significance of the use of “sneer” quotes around the word ‘true’, here, see Stove (2011).

authorized person A_1 recorded: ‘physician B obtained a blood sample from patient C’

authorized person A_2 recorded: ‘physician B did not obtain a blood sample from patient C’.

This means, however, that the extension of the RIM’s Act — meaning the totality of acts in reality described through Acts — overlaps only partially with that of the BFO category of occurrents. This is because Act, for the RIM comprehends not only the records which result from documenting processes *that have happened*, but also counterpart records of processes that *did not happen*, that *can happen*, *are intended to happen*, *are requested to happen*, and so forth, in reflection of the RIM’s distinguished mood codes. Or as HL7 itself formulates the matter: ‘The moodCode distinguishes among Acts that are meant as factual records, records of intended or ordered services, and other modalities in which acts can be recorded.’¹ Acts will thus include also, for instance, (records of) intended but cancelled surgeries. The fact that there can be Acts, in the sense of the RIM, which never happen is from this point of view understandable. Unfortunately however this outcome is in conflict with the RIM’s own description of Act (see Table 1, again) as representing *the actions that are executed*.

8. Conclusion

Our purpose has been to show how the RIM’s backbone classes line up, or rather fail to line up, with the fundamental categories of the BFO ontology. If BFO, as we believe, comes close to capturing categorical distinctions at the heart of common sense, then this failure implies major shortcomings in the RIM — shortcomings for example as concerns learnability, coherence, comprehensiveness, and stable evolution. As we saw, the HL7 has presumed that the problems of dialect formation which had plagued v2 would be resolved once the RIM was brought into play. On the basis of this presumption the HL7 organization has gone on to promulgate manifold varieties of RIM-conformant ‘v3’ standards, and these new standard — even in spite of their manifest shortcomings — have been incorporated into multiple health information technology

¹ Health informatics: – HL7 version 3 – Reference information model – Release 4 (Document: ISO/HL7 21731:2011(E)), <http://standardsproposals.bsigroup.com/Home/getPDF/1361>, last accessed November 28, 2012.

initiatives, where they have sometimes been associated with conspicuous failures, some of them on a national scale.¹ On the other side, however, there are also positive signs which are slowly but surely beginning to manifest themselves, and we are pleased to observe that the claims made on behalf of the RIM are increasingly being treated with suspicion in healthcare informatics circles.²

References

- Austin, J. L. (1962). *How to do Things with Words*, Oxford: Oxford University Press.
- Grenon, P., Smith, B., and Goldberg, L. (2005). “Biodynamic Ontology: Applying BFO in the Biomedical Domain”, in D. M. Pisanelli (ed.), *Ontologies in Medicine*, Amsterdam: IOS Press, 20–38.
- Grieve, G. (2012). “HL7 needs a fresh look because V3 has failed”, <http://www.healthintersections.com.au/?p=476>, last accessed November 12, 2012.
- Rector, A., Nolan, W., and Kay, S. (1991). “Foundations for an electronic medical record”, *Methods of Information in Medicine*, 30: 179–86.
- Searle, J. (1969). *Speech Acts. An Essay in the Philosophy of Language*, Cambridge: Cambridge University Press.
- Smith, B. (1990). “Towards a History of Speech Act Theory”, in A. Burkhardt (ed.), *Speech Acts, Meanings and Intentions. Critical Approaches to the Philosophy of John R. Searle*, Berlin/New York: de Gruyter, 29–61.
- Smith, B. (2003). “John Searle: From Speech Acts to Social Reality”, in B. Smith (ed.), *John Searle*, Cambridge: Cambridge University Press, 1–33.
- Smith, B. and Ceusters, W. (2006). “HL7 RIM: An Incoherent Standard”, *Studies in Health Technology and Informatics*, 124: 133–138.

¹ <http://hl7-watch.blogspot.com/2011/08/uk-scrapping-national-health-it-network.html>, last accessed November 12, 2012.

² Grahame Grieve, “HL7 needs a fresh look because V3 has failed”, <http://www.healthintersections.com.au/?p=476>, last accessed November 12, 2012.

- Smith, B. and Ceusters, W. (2010). "Ontological Realism as a Methodology for Coordinated Evolution of Scientific Ontologies", *Applied Ontology*, 5, 139–188.
- Stove, D. C. (2011). *Scientific Irrationalism: Origins of a Postmodern Cult*, New Brunswick: Transaction.
- von Wright, G. H. (1963). *Norm and Action: A Logical Inquiry*, London: Routledge & Kegan Paul.