

**Foundation for the Electronic Health Record:
An Ontological Analysis of the HL7's Reference Information Model**

Lowell Vizenor^{1,2}, Barry Smith^{1,2} and Werner Ceusters³

1. Institute of Formal Ontology and Information Science, University of Saarbrücken,
Germany
2. University at Buffalo, Department of Philosophy, Buffalo, New York, USA
3. European Institute for Ontological Research, Saarbrücken, Germany

Summary

Despite the recent advances in information and communication technology that have increased our ability to store and circulate information, the task remains of ensuring that the right sorts of information reach the right sorts of people. In what follows we defend the thesis that efforts to develop efficient means for sharing information across healthcare systems and organizations would benefit from a careful analysis of human action in healthcare organizations, and that the communication of healthcare information and knowledge needs to rest on a sound ontology of social interaction. We illustrate this thesis in relation to the HL7 RIM, which is one centrally important tool for communication in the healthcare domain.

Keywords: Ontology, Speech Acts, HL7, Electronic Health Record, Organizations

1. Introduction

Information and communication technology has not only altered the way that medical information is stored, shared, and generated across and within healthcare organizations; it also promises to increase the efficiency and cost-effectiveness of healthcare. The hope is that the electronic health record (EHR) will facilitate the diffusion and dissemination of healthcare information through the use of clinical guidelines, protocols and messaging standards. Still, many difficulties must be overcome before a comprehensive EHR is realized. One such difficulty arises from the fact that the single doctor-single patient nexus has been largely superseded by a regime in which the typical patient is managed by a team of health care professionals, each specializing in one aspect of care.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 – for example, the immunology and the oncology departments – have different disciplinary cultures and employ different terminologies to talk about what are putatively the same phenomena. We argue that 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 would benefit from a careful analysis of human action that is based on a sound ontology of social reality. In order to demonstrate this thesis we turn our attention to the work of the Health Level 7 (HL7), one of several ANSI-accredited Standards Developing Organizations operating in the healthcare arena [1].

HL7 has advanced a widely used messaging standard – one which has recently been established as mandatory for communication between US Federal Government-funded healthcare organizations – that enables healthcare applications to exchange clinical and administrative data in digital form. HL7 seeks 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. The breadth of domains within the remit of HL7 has inspired the development of a powerful abstract model of patient care called the Reference Information Model (RIM) [2], which is intended to serve as a unified framework for the sharing of information and the usage of data across different healthcare domains. Moreover, the RIM can be viewed as an ontology insofar as it amounts to a representation of the healthcare domain that is then mapped by the allowed bits of HL7 syntax.

As we shall see, the RIM in its documentation explicitly acknowledges an indebtedness to philosophy. On the other hand, however, it does not situate its own core theory of acts within a sound philosophical ontology. We accordingly set the scene with a discussion of philosophical ontology and with arguments to the effect that sound ontology should serve as a foundation for a comprehensive EHR and associated messaging systems. We accomplish this through an examination of the HL7 RIM in which we attempt to determine its expressive adequacy from an ontological point of view. Familiarly, the authors of the RIM adopt an explicitly act-centered view of healthcare [3], and they draw on the theory of speech acts as developed by J. L. Austin in his *How to Do Things with Words* [4] for their underlying framework. The theory of speech acts is part of a more comprehensive ontology of organizations, and in order to establish how far the HL7 RIM succeeds in providing a framework that does justice to the way speech act phenomena actually work we need to establish the degree to which it is compatible with this more comprehensive ontology. To this end we pay attention to the sorts of preconditions that must be satisfied for the successful performance of a speech act and to the sorts of entities to which speech acts give rise (obligations, claims, commitments, and so forth).

2. Ontology and the HL7 Reference Information Model

2.1 The Act-Centered View of Healthcare

The RIM's act-centered view of healthcare is based on the assumption that any profession or business, including healthcare, consists primarily of a series of intentional actions on the part of responsible actors. The varieties of such actions include: actions of clinical observation; actions of assessment of health conditions such as the taking of diagnoses; actions of providing treatment services such as surgery and physical therapy; actions of assisting, monitoring, attending and training; actions of administering education services to patients and their next of kin; actions of providing notary services

such as the preparation of an advanced directive or a living will; actions of editing and maintaining documents, and so forth.

A special sort of intentional actions are those linguistic actions known as *speech acts*. J. L. Austin [4] and John R. Searle [5] are normally taken to have been the first to emphasize that what we can do with words goes well beyond uses of language of the statement-making sort. We can make requests, ask questions, give orders, make promises, and so on, all of which are kinds of actions marked by the fact that their very utterance brings about some result. The authors of the RIM recognize the importance of paying special attention to the ways that speech acts such as orders and requests are a kind of doing that goes beyond the simple reporting of facts. In addition, they recognize also that these acts stand in systematic relations to the other acts occurring within the framework of a healthcare organization and that the more accurately a healthcare information model captures these systematic relations between speech acts and other kinds of action the more information will it make available to its users.

But of course there are many features of healthcare that go beyond the category of action. These include the *participants* of the actions themselves, both agent and patient; they include the *roles* these participants play in actions, their *authority* to perform a given action, and they include the sorts of entities to which these actions give rise such as *obligations* and *claims*. They also include physical objects such as buildings, and healthcare contexts/institutions such as wards and laboratories. They include substances such as eye drops and blood samples, and events such as births and deaths. 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 these entities and their proper relations to each other.

The authors of the RIM respond to this need by organizing the relevant healthcare information along the lines of what they call the six “backbone” classes of the RIM which they call: Act, Entity, Role, Participation, ActRelationship, and RoleLink [3]. In what follows we discuss these classes in relation to the fundamental ontological categories distinguished by philosophers in dealing with speech acts and similar phenomena.

For the purposes of clarity, we capitalize the first letter of HL7 terms. When terms such as ‘entity’, ‘act’, ‘action’, and ‘role’ appear without initial capitals, then their common meaning is intended.

2.2 Continuants

All real world entities of the sort which we encounter in the healthcare domain fall into one of two exclusive categories of continuant and occurrent [6].

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 up into the two sub-categories of *independent* and *dependent*. Typical examples of independent continuants from the healthcare domain include physicians, administrators, patients and family members, medical supplies, and

records and documents in filing cabinets. Each of these entities continues to exist through time even as it undergoes changes in personnel or changes in the organization of wards and departments. Examples of dependent continuants are the *states*, *properties*, *qualities*, and *roles* of patients, administrators and so forth.

The *EntityClass* represents the closest analogue in the RIM to what we have just called an independent continuant. This class comprehends all those sorts of physical things or groups of physical things which can participate in an action as perpetrator, target or beneficiary such as people, places, organizations, medical tools, supplies, and so on. The *EntityClass* does not include those actions and events in which Entities participate – which we shall classify as occurrents – and in this respect the analogy between Entities and continuants holds.

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 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* of healthcare personnel such as the ability to speak Spanish or to perform complex medical procedures; the *roles* that participants play in actions; and their *authority* to perform a given action. 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 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 care giver to perform certain activities that fall under the jurisdiction of the medical authority in question [7]. The *LicensedEntityRole* is an HL7 Role that resembles what we have called a relation of authority, which is a dependent continuant. 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 [7]. HL7 Roles should not be confused, though, with what we have called roles or what in the linguistic literature about semantics is known as *semantic role*, *semantic case*, *case role*, *thematic role*, *theta role* (in generative grammar), and *deep case* (in case grammar) [8]. Thus, while there are some similarities between HL7 Roles and dependent continuants, there are also many significant differences, to one of which we now turn.

2.3 Universals and Particulars

One of the specializations of RoleClass is *RoleClassOntological*. This class includes such Roles as *has generalization*, *instance* and *subsumed by*. This terminology might be found problematic if it is not recognized that from an ontological point of view it is important to keep straight two distinctions. On the one hand is the distinction between a role and the individual bearer upon whom that role is existentially dependent. On the other hand is the distinction between a universal and the individual or individuals by which it is instantiated [9]. The distinction between particulars and universals is important since it is precisely universals which provide the basis for the classifications of particulars. The universal is, thus, distinct from its extension, the set of its actual instances at any given time. Universals exist when and where they are instantiated. In turn, every particular is an instance of one or more universals. In contrast to particulars, universals can exist at different places simultaneously.

As particulars correspond to proper names and other singular referring expressions, universals correspond to general terms. They are that in reality which their particular instances, both actual and possible, share in common. Both dependent and independent continuants instantiate corresponding universals. Just as humans, hospitals, kidneys, and so on, instantiate universals, so too do dependent continuants such as their roles, states and qualities. Just as it is common in a healthcare setting to find multiple instantiations of the universal *human*, so too it is common to find multiple instantiations of the same role. For example, the role *nurse practitioner* is multiply instantiated whenever a hospital has more than one nurse practitioner; but in each (particular) case it is the same role (universal) that is instantiated.

All of this means that the RoleClass of the HL7 RIM is ontologically heterogeneous; that is, it comprehends items from different ontological categories. Some HL7 Roles (and RoleLinks) represent dependent continuants such as relations of authority, others represent relations between universals and their instances such as *has generalization of*.

2.4 Occurrents

Occurrents (also called events, processes, activities) are defined as being such that they unfold themselves in their successive phases. Thus in contrast to continuants they never exist in full in any single instant of time. Some occurrents, such as surgical operations, have a beginning, a middle, and an end; others, such as inventorying, are cyclical. The life of a patient is an occurrent, as also is the 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 the plan. (Plans themselves however are continuants: thus they endure continuously through time until they reach the point of complete execution.)

2.4.1 ActClass

The *ActClass* represents the closest analogue in the RIM to occurrents, but it is at best a weak analogy. For the ActClass does not comprehend actions themselves but rather only

the records which arise when such actions are documented by a healthcare professional in either a clinical or administrative context. Moreover the notion of *record* is itself hereby an abstraction. For it does not correspond to actual documents but is rather an encapsulation of specific sorts of information pertaining to the actions themselves (e.g. to the nature of the action performed and to the time when it was executed) and also to the recording of the actions (e.g. when and in what language it was recorded). Each record of an action in turn is referred to in the RIM as an Act-instance.

2.4.2 Primary and Secondary Acts

In order to make these matters somewhat clearer, we introduce a distinction between a *primary* act and a *secondary* act. The primary act of every Act-instance is an act of documenting some other secondary act. Consider the case where physician A documents that physician B obtained a blood sample from patient C. In this case, the primary act is the act of documenting on the part of physician A and the secondary act is the obtaining of blood on the part of physician B. The person responsible for the Act-instance is the one who documents the action and not the one who performs the real world action (i.e., obtaining a blood sample). In other cases, though, the person who documents the action of obtaining a blood sample is the same person who obtains the blood sample. In this case the same person is responsible for the real world action and its documentation [3]. Since every Act-instance is of the same kind – an act of documenting – the primary act cannot serve as an informative basis for classifying Act-instances. Thus the RIM must rely on the secondary acts to provide the basis for the classification of Act-instances.

It is important to keep in mind the distinction between the primary and the secondary act, if one wants to understand the rationale behind the RIM's design. Following Rector and Nolan, the authors of the RIM view *faithfulness* to the clinical history and care of the patient as the fundamental criterion for what gets included in the medical record:

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 [10].

Every Act-instance is committed to accurately and truthfully representing who documented what and when and where it was documented, but is not likewise committed to the truth of the content of the secondary act. Thus 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 HL7, is whether or not the proposition “physician A 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 through the descriptions of secondary acts, but since these descriptions are always attributed to someone, it is possible to have two Act-instances which describe the same real world activity yet disagree; and because each is properly attributed to its author, such disagreements can exist side by side and the task of determining which is true be left to the recipient of these Act-instances [3].

2.4.3 Moods and ActRelationships

The ActClass does not however comprehend only records of events that have happened, but also of those that can happen, are happening, are intended to happen, or are requested/demanded to happen. The authors of the RIM distinguish to this end between two aspects of an Act-instance: its mood and its content. Consider the following examples: physician A *performed* a laparoscopic cholecystectomy, physician A is *performing* a laparoscopic cholecystectomy, and physician A *intends to perform* a laparoscopic cholecystectomy. All three examples are identical with respect to their content, i.e. they make reference to performing a laparoscopic cholecystectomy, but they differ with respect to their mood: the first is something that has happened, the second something that is happening and the third something that will happen. Moods, then, are intended to capture the way healthcare activities progress in the real world, from the initial decision to perform the action, through the planning and ordering of the action to the completion of the action. The different Act-instances of a laparoscopic cholecystectomy (e.g. defined, planned, ordered, and executed) are linked together in the RIM by ActRelationships. Important here is the *ActRelationshipSequel*. This class is intended to capture the relations between, say, the issuing of an order for a blood sample and the obtaining of a blood sample. In this case, the relation between these two Act-instances would be that of fulfillment. Other examples of ActRelationships in the RIM include such things as *have component*, *has goal*, *reason*, and so on [7].

Finally, we turn to the ParticipationClass, which associates Entities in a Role with the specific way in which they are involved in an Act-instance. For instance, a healthcare provider can participate in an action either as performer, recipient, or beneficiary. Other ways in which an Entity may participate in an action are as *data entry person*, *donor*, *legal authenticator*, *consultant*, and so on [7].

2.4.4 Discussion

It will by now be clear that the ActClass only partially overlaps with the category of occurments. First, with the possible exception of the act of documenting, the ActClass comprehends occurments (actions) only indirectly, via their records, and records – even the abstract records conceived of by the RIM – are themselves occurments. Moreover, there are many occurments common to healthcare settings which are not represented at all. No reference is made, for instance, to natural events such as a patient's heartbeat. If a given patient's heartbeat falls within the perspective of the RIM at all, then this can be so only because there is some corresponding record of an action that is attributed to someone. Thus a patient's heartbeat may be recorded as observed, but then there is no record of the heartbeat itself; rather the record is of the observation by, say, a physician.

In addition, many of the specific items that fall under the ActClass are not occurments. The RIM classifies clinical documents in the ActClass [7]. And in its description it lists several characteristics of a clinical document: (1) persistence – a clinical document continues to exist in an unaltered state, for a time period defined by local and regulatory requirements; (2) Stewardship – a clinical document is maintained by a person

or organization entrusted with its care. Yet, persistence, as we have seen, is a characteristic not of acts but of continuants. A clinical document exists *in total* at a given time, it does not unfold through successive stages over time. Second, a person cannot *maintain* a collection of attributed acts in the way in which he might maintain a document or collection of documents. Similarly, the RIM locates *financial account* in the ActClass. This in turn collapses the distinction between the account itself, which is a dependent continuant, the documents which record it, which are independent continuants, and the associated actions including (i) those that bring the account into existence, (ii) those that modify the account, and (iii) those that nullify it or bring it to an end.

Finally, of those Acts that qualify as genuine actions only a few – as stated in the HL7 documentation – count as speech acts (e.g. *consent*, *inform* and *financial transaction*). The failure to distinguish between acts (occurrents) on the one hand and documents (continuants) on the other reveals that the RIM does not track the distinction between continuants and occurrents. As a result information that comes with the tracking of this distinction is lost; moreover, it can be anticipated that those involved in coding on the basis of the RIM will confront additional obstacles in virtue of the ontologically counter-intuitive character of its six backbone classes.

2.5 Summary on Ontology and the RIM

Our purpose thus far has been only to show how these backbone classes line up with the fundamental categories of ontology. From an ontological point of view, the RIM has two major shortcomings. First, it is not comprehensive. Many items fundamental to healthcare – especially natural events – do not find a home in the RIM. Second, the RIM maintains no clear distinction between roles *qua* dependent particulars and roles *qua* universals. Third, the category of dependent continuants is not systematically tracked.

In what follows, we look at the role that speech acts play in the RIM, and then show how the categories of ontology provide a framework for systematically tracking the sorts of actions and states that are common to healthcare.

3 Speech Acts

3.1 J. L. Austin and Performatives

In *How to Do Things with Words* Austin sketched an account of performative utterances, which he defined as those uses of language which are themselves a kind of action and whose very utterance brings about some result. He noted that speech acts often involve a certain ritual aspect. Of an utterance like ‘I promise to mow your lawn’ we ask not whether it is true, but whether it is successful. The conditions of success for performatives Austin called *felicity conditions* and he saw them as ranging from the highly formal (such as those governing a judge pronouncing a sentence), to the informal conventions governing expressions of gratitude or sympathy in the circumstances of everyday life. By the end of *How to Do things with Words* Austin has given up on the idea of a theory of performatives as such. This is because he has reached the conclusion

that all utterances are in any case performative in nature, and thus he replaces his earlier goal of a theory of performatives with the goal of a theory of speech acts in general [11].

3.2 Illocutionary Force and Propositional Content

To see how such a theory would look we need to turn to Austin's successor Searle. [5] Searle points out that in each speech act we can distinguish abstractly two components: the *type or quality* of the act (sometimes called its illocutionary force) and the (normally propositional) *content* of the act. Each can vary while the other remains constant, as we can demand or request or express our desire *that a blood test be performed* or request *that more medical supplies be ordered* or request *that more medical personnel be staffed*, and so on. The following scheme represents this distinction, whose similarity with the RIM's mood/content distinction will be evident:

S(P)

Here S represents the illocutionary force of the speech act and P represents the propositional content. With this distinction in mind we now turn to classification of speech acts [5].

3.3 Directions of Fit

In his "A taxonomy of Illocutionary Acts", [12] Searle introduces the notion of a "direction of fit." The idea is that the direction of fit between the propositional content of a speech act and that in the world to which it relates differs with the different sorts of speech acts. Some speech acts – what Searle calls assertives, including statements, assertions, and descriptions – are supposed to match up to an independent reality, and to the extent that they do so we may say that they are true or false. These speech acts have a *word to world* direction of fit.

For other speech acts the direction of fit is from *world to word*. Members of the directive class—orders, commands, requests, etc.—and the members of the commissive class—promises, vows, pledges, etc.—are not supposed to match an independently existing reality. Rather, they are supposed to bring about changes in the world so that the world matches the propositional content of the speech act.

4 An Ontological Analysis of Speech Acts

Austin and Searle focus their attentions on speech acts as part of their work in the philosophy of language. The needed ontological backbone of their theory was however provided fifty years earlier by the German philosopher Adolf Reinach in his [13], which has strong claims to be the work in which the theory of speech acts was first set forth in a systematic way. [14]

4.1 The Promise

Reinach treats what he called the *theory of social acts* as a branch of ontology – it is part of a general ontological theory of the structures of social interaction. For the purposes of our ontological analysis of speech acts we will take inspiration from Reinach and begin with a discussion of one specific sort of speech act, namely the promise.

Speech acts are social acts and as such they involve at least two people. Every promise requires a *promiser* and a *promisee*. In this sense, speech acts are opposed to solitary actions (e.g. the rehearsal of a difficult medical procedure in one's head). In the case of the promise, there exists a dependence relation between the promise and promiser on the one hand and a dependence relation between the promise and the promisee on the other hand. Moreover, the dependence is in each case *one-sided*: the promiser and promisee are persons; they themselves do not require the existence of the promise in order to exist [6]. What we can say, however, is that the *roles* those persons play within the promise-making context and the promise itself are mutually dependent on one another. Two or more objects are mutually dependent if each can exist only in combination with the others. At one level, these relations represent relations among individual instances: Mary makes this promise to John at this particular point in time and in this particular context. But they can also be seen as representing relations among the corresponding universals or general classes which such individuals instantiate—here, among the classes: *promiser*, *promise*, *promisee*, etc.

4.2 Satisfaction Conditions

To do justice to the ontological structure of the promise we must now recognize certain further elements that must be present in order for there to be a successful issuance of a promise:

- (1) An act of speaking or more generally of message-sending (on the part of the promiser)
- (2) An intention (the promiser must intend to perform some action)
- (3) An act of uptake, acceptance, or acknowledgement on the part of the promisee
- (4) A positive evaluation in relation to the given action (the promisee must, for example, desire that the promiser perform F).

Thus (1), the promiser must utter something to the effect that 'I promise to do F' and this utterance must be accompanied (2) by a corresponding intention to perform this action (and to do so in order to fulfill the promise). This is needed in order to ensure that (3) the promise can be apprehended by the promisee – and this is important for no other reason than that the promisee must have the opportunity to decline or accept the promise. Finally, only if the promisee accepts (4) is the ontological structure of the promise realized.

In addition to these conditions there are some additional background conditions, for example, that in some cases the promiser must have the *authority* to promise to do F and/or to perform the corresponding action, or that the promisee must have the authority to accept or decline F. Only some healthcare personnel have the authority to promise to a patient that he will receive a certain medical procedure. Only under certain circumstances does the guardian of a patient have the authority to grant permission on the patient's

behalf for a medical procedure or treatment. Similarly, we can point to a condition to the effect that if the structure of the promise is to be realized then there should exist no prior commitments on the part of the promiser or the promisee that would *vitiate* the promise. For example, a prior obligation to perform a kidney transplant can vitiate a subsequent promise to perform another kidney transplant when the same kidney is involved. Such conflicts may exist – and may be recognized by a system for keeping track of information within a healthcare organization – even when there is no record of a corresponding observation that is attributable to someone. Finally, we can state some general background conditions: promiser and promisee should share a common language, and there must exist a general background of trust if the entire social practice of promise making itself is going to work.

4.4 Successor States

But we are not yet done. If all the above conditions are satisfied, then two new continuant entities – entities of a normative type – come into existence. For a promise to do F gives rise on the one hand to an *obligation* to do F (on the part of the promiser) and on the other hand to a *claim* that F be done (on the part of the promisee). These we can call the *successor states* of a promise. The promise itself is composed of a three-part structure: the act of speaking, the act of uptake, and the content (F). The promiser and promisee, however, as well as the obligation and claim are not component parts of the promise; rather they do belong to the entire surrounding structure.

It will be clear that to do justice ontologically to a speech act such as promising we also need to take explicit account of the dimension of time. The act of promising is in the simple case an event in time (an occurrent) that exists only so long as it takes for the promiser to utter the statement and for the promisee to register it. (More complex cases arise, for example, when a promise is made via email.) Most hospital information systems now use Patient Order Entry Systems [15], meaning that there is no direct contact between the person who orders and the person who has to execute the order.

Obligation and claim, in contrast, are not occurrent *events* but rather continuant *states*. This means that they exist unchangingly from the time of the successful issuance of the promise to the time when the content of the promise is fulfilled or satisfied. Promises, it should be noted, do not *cause* their own fulfillment. Rather, the relation involved is something weaker than causation. Following Reinach, we can talk here in terms of a tendency to be realized. Just as genes have a tendency to be expressed in the form of proteins, and just as bodies have a tendency to fall when dropped, so promises have a tendency to be fulfilled. In all such cases mentioned, the tendencies in question can be blocked.

5 The Order

While Reinach focuses primarily on the promise he distinguishes a range of other varieties of speech acts, and sketches for some of these the beginnings of an ontological

analysis. Among the cases dealt with are: accepting, advising, agreeing, allowing, answering, apologizing, asking, commanding, complaining, confessing, denying, informing, insulting, praising, praying, requesting, taking an oath, and thanking.

For HL7 purposes it is the case of *ordering* which is most important. HL7 does not have a class devoted specifically to orders themselves, and the HL7 documentation, for example in the description of the ActClassSupply, tells us only that ‘Supply orders and deliveries are simple Acts that focus on the delivered product.’ However we surmise that the term ‘order’ is best understood in the HL7 framework as a speech act in which one participant with a corresponding authority requests another participant to do something, for example to deliver certain supplies or to perform a certain blood test.

Each order thus involves at least two people: an orderer and an orderee. (We ignore the complications which arise in those cases where the orderee is a plural subject, for example a team, or where it is another healthcare institution.) In this sense, speech acts are opposed to solitary actions (e.g. the rehearsal of a difficult medical procedure in one’s head). There then exist dependence relations between order and orderer and between order and orderee exactly analogous to those existing in the promising case. Here, too, we can parse the ontological structures involved in such a way as to distinguish two mutually dependent *roles* the separate persons play within the order-giving context.

The following further elements must also be present in order for there to be a successful issuance of an order:

- (1) An act of speaking or more generally of message-sending (on the part of the orderer)
 - (2) An intention (the orderer must (a) intend that the orderee perform some action and in such a way that (b) the orderee performs this action in virtue of the giving of this order)
 - (3) An act of uptake (registering, understanding, accepting) on the part of the orderee
 - (4) A relation of authority, in virtue of which the orderer is entitled (a) to give an order to perform this action that is (b) directed to the orderee.
- (3) will be combined with the establishment of a tendency to perform the given action. 2(b) ensures that the act of ordering is different from merely wishing (which can be fulfilled even if an action is performed completely independently of any act of communication). 4(b) ensures that the order is distinct from a mere request, where an analogue of the relation of authority is absent.

Here again there are also successor states. The orderee acquires an *obligation* as a result of the order to carry out the corresponding action, which is once again a continuant state which exists unchangingly from the time of the successful issuance of the order to the time when its content is fulfilled through the corresponding action on the part of the orderee. And here again the relation linking the issuance of the order to the performance of this action is something like a tendency, a relation weaker than causation.

6 Conclusion

The HL7 RIM, although originally designed to assist the difficult process of writing HL7 messages, is gaining importance as an information model of healthcare delivery in its own right. This should come as no surprise, since the development of the RIM was initiated at a time when information modeling was considered to be a key area of activity in healthcare standardization bodies all over the world. In a historical perspective, the first messages were rather simple: a laboratory information system wanted to receive some administrative data from a patient stored in the ADT system. Later, results from the lab had to be displayed on point of care systems. Soon, order entry became important, not just for lab systems, but also for any other department delivering technical services such as radiology or neurophysiology. Nowadays, departmental systems are all over the place inside a hospital and a wide range of clinical data is registered and exchanged over many systems, which may even extend across multiple hospitals.

The way information systems are set up in a hospital, and in particular the flow of communication amongst them, is a near replica of how the hospital itself is organized. The major advantage of the HL7 RIM is that it recognizes the major players inside such an organization together with the activities they engage in.

In this paper we analyzed the HL7 RIM from the perspective of speech act theory and submitted its various classes to ontological analysis. We discovered that the RIM is marked by a number of problems, above all when it comes to taking dependent continuants properly into account. Although the RIM intends to represent registrations of relevant medical acts in a way which acknowledges the distinction between what we have called primary and secondary acts. Yet this very distinction is not adequately maintained in all the parts of the RIM. It is true that this shortfall may be of little practical importance for the purposes of messaging; but if the RIM has the goal of being used as a reliable and effective ontology of medical acts in the future, then the situation would be quite different. For delivering adequate and cost-effective care, it is mandatory that one can trust an ordered test not just to be carried out, but also to be registered both as having been ordered (to avoid it being ordered a second time for no other reason than that it was not known to have been ordered before) and as having been carried out. The way this and related types of information are coded in HL7 messages may be sufficient for interpretation by HL7-trained humans, but it does not allow algorithms to grasp the underlying differences. We propose that to make this possible in the future correction of the RIM should be undertaken on the basis of a more fine-grained ontological analysis along the lines set forth above.

Acknowledgements

This work was supported by the Alexander von Humboldt Foundation under the auspices of its Wolfgang Paul Program and by the EU FP6 Network of Excellence on Medical Informatics and Semantic Data Mining.

References

- [1] http://www.hl7.org/library/standards_non1.htm
- [2] http://www.hl7.org/library/data-model/RIM/modelpage_mem.htm
- [3] <http://www.hl7.org/Library/data-model/RIM/C30202/rim.htm>
- [4] J. L. Austin. *How to do Things with Words*, Oxford University Press, Oxford, 1962.
- [5] J. Searle, *Speech Acts. An Essay in the Philosophy of Language*, Cambridge University Press, Cambridge, 1969.
- [6] Grenon, Pierre, Barry Smith, Louis Goldberg: *Biodynamic Ontology: Applying BFO in the Biomedical Domain*. in: Pisanelli, Domenico M. (ed.): *Ontologies in Medicine: Proceedings of the Workshop on Medical Ontologies*, Rome, October 2003. IOS Press, Amsterdam.
- [7] <http://www.hl7.org/Library/data-model/RIM/C30202/vocabulary.htm>
- [8] T. Payne, *Describing morphosyntax: A guide for field linguists*, Cambridge University Press, Cambridge; New York, 1997.
- [9] B. Smith, *On substances, accidents and universals: in defense of a constituent ontology*. *Philosophical Papers*, 26 (1997), 105–127.
- [10] A. Rector, W. Nolan and S. Kay, *Foundations for an electronic medical record*. *Methods Inf Med*, 30 (1991): 179-86.
- [11] B. Smith, *From speech acts to social reality*. in: B. Smith (Ed) John Searle. Cambridge University Press, Cambridge, 2003.
- [12] J. Searle, *A taxonomy of illocutionary acts*, in: Keith Gunderson (Ed.) *Language, Mind and Knowledge*, Minnesota Studies in the Philosophy of Science, Vol. VII, University of Minnesota Press, Minneapolis, 1975, pp 344-367.
- [13] A. Reinach, *Die apriorischen Grundlegen des bürgerlichen Rechts*, *Jahrbuch für Philosophie und phänomenologische Forshung*, I/2, 685-857, English translation as *The Apriori Foundations of the Civil Law*. By J. F. Crosby in *Aletheia*, 3 (1983), 1-142.
- [14] K. Mulligan, *Promisings and other social acts: their constituents and structure*, in: Mulligan (Ed), *Speech Act and Sachverhalt*. Reinach and the Foundations of Realist Phenomenology, Nijhoff, The Hague, pp.29-90.
- [15] <http://www.ncvhs.hhs.gov/010201h6b.htm>