Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Biomedical Informatics xxx (2005) xxx–xxx
Journal of Biomedical Informatics
www.elsevier.com/locate/yjbin

Strategies for referent tracking in electronic health records

Werner Ceusters [a,c,*], Barry Smith [b,c]

[a] European Centre for Ontological Research, Univerisität des Saarlandes, Postfach 151150, D-66041 Saarbrücken, Germany
[b] Institute for Formal Ontology and Medical Information Science, Saarbrücken, Germany
[c] Department of Philosophy, University at Buffalo, NY, USA

Received 14 June 2005

Abstract

The goal of referent tracking is to create an ever-growing pool of data relating to the entities existing in concrete spatiotemporal reality. In the context of Electronic Healthcare Records (EHRs) the relevant concrete entities are not only particular patients but also their parts, diseases, therapies, lesions, and so forth, insofar as these are salient to diagnosis and treatment. Within a referent tracking system, all such entities are referred to directly and explicitly, something which cannot be achieved when familiar concept-based systems are used in what is called "clinical coding." In this paper, we describe the components of a referent tracking system in an informal way and we outline the procedures that would have to be followed by healthcare personnel in using such a system. We argue that the referent tracking paradigm can be introduced with only minor—though nevertheless ontologically important—technical changes to existing EHR infrastructures, but that it will require a radically different mindset on the part of those involved in clinical coding and terminology development from that which has prevailed hitherto.

Keywords: Referent tracking; Electronic Health Records; Ontology; Biomedical terminology; Semantic interoperability

1. Introduction

Electronic health records (EHRs) consist primarily of descriptions of a patient's medical condition, the treatments administered, and the outcomes obtained. These descriptions are about concrete entities in reality: for example about the particular pain that the particular patient John experienced in his chest on this specific day; or about the particular pacemaker—with its specific serial number assigned to it by its manufacturer—that was implanted in John during the particular surgical procedure that started at a precise moment in time on a certain day.

The descriptions contained in current EHRs contain very few explicit references to such entities. This lack of explicit reference is usually a minor problem for human interpreters, but it makes an accurate understanding of EHR data nearly impossible for machines. This is because reference resolution in running text (still the most common format for descriptions in EHRs) is one of the hardest problems in natural language understanding [1]. But even those EHR systems which incorporate data in more structured formats, for example by resorting to controlled vocabularies, terminologies, or ontologies, are in no better shape in this respect. This is because the terms or codes contained in the latter are used simply as an alternative to what would otherwise have been registered by means of general terms in natural language. By picking a code from such a system and then registering that code in an EHR, one refers generically to *some* instance of the class represented by the code. It is still left at best only partially, and indirectly, specified which particular instance is intended in concrete reality.

This has some obvious consequences. When a patient suffers from the same type of disease and exhibits the same kinds of symptoms on two successive occasions, then the descriptions of these conditions using codes from a terminology will be identical. When another patient suffers from

* Corresponding author. Fax: +49 0 681 302 64772.
E-mail address: Werner.Ceusters@ecor.uni-saarland.de (W. Ceusters).

1532-0464/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.jbi.2005.08.002

the same type of disease and exhibits similar symptoms in his turn, then the resulting descriptions will also be identical to those relating to the first patient.

Certainly some of the associated references to specific persons, dates and places will be different; but the information conveyed by such references is not always sufficient to assess whether the same (i.e. numerically identical) or different (but qualitatively similar) entities are being referred to. Is the closed atlas fracture referred to by using the SNOMED-CT concept code 269063003 in John's EHR at time t1 the very same fracture as that which is referred to by means of the same code at time t2? If t1 and t2 are not far apart in time, then it is probable that the answer to this question is 'yes;' if there is a large gap, then poor John probably broke his neck twice. But who will be able to tell the difference on the basis of what is contained in the corresponding record? One can assume that, under normal circumstances, a fracture will not have healed within a time frame of 2 weeks, and also that it will have healed after 2 months. But there is an intermediate period during which such assertions cannot be safely made. And what if it is not a fracture that is being referred to, but a disease such as diabetes, or a malformation such as a webbed neck? In such cases, as is obvious (again) to human beings but not to the machine, the very same disorder or malformation may be referred to at successive times and in a plurality of ways over the whole of the patient's lifetime.

Of course a given record may contain descriptions about when a specific disease was considered to have been cured (for example a wound to have been healed), so that, if a new entry about the same type of disease is made thereafter, then it will be possible to infer that it pertains to a new instance. Unfortunately, however, patients tend to visit their physicians primarily when they are ill, and not when they are cured, so that EHRs contain far more traces of information pertaining to the initial phases of a disease than to its successful treatment. And even where information pertaining to the termination of a disease has been registered, it will still often be very hard for software applications to understand this information properly and to make the appropriate inferences.

Similar considerations can be made regarding descriptions in the EHRs of two different patients. Note, first of all, that one cannot assume that if the same code is used in two such records then they refer to two distinct entities. Temperatures and disorders depend uniquely on their particular bearers, and so they necessarily constitute different entities when referred to by means of corresponding codes or general terms in the EHRs of different patients. But SNOMED-CT has codes for places such as "swimming pool," or relatives such as "father." Obviously, two different patients may share a common father, and they may have visited the same swimming pool from which they then obtained (different) nasty viral warts. They may or may not have received transfusions of blood from the same donor, thereby becoming HIV-positive, and so forth. In these cases it is important, at least from an epidemiological perspective, that inferences can be reliably made as to whether it is the same or different entities in reality which are being referred to.

Only a few codes in systems such as SNOMED-CT refer always to the same particular entity. This is so for instance in the case of codes referring to countries such as *Belgium*, the *United States of America*, and so forth, or to specific protocols, social security plans or named drug regimes. But this facility, which is a strength from the point of view to be defended in this paper, is in fact to be assessed as an ontological weakness from the perspective of SNOMED-CT itself. This is because the latter system—like so many others—is not able to distinguish between terms, such as "*country*," used in a general sense, and terms, such as "Belgium," that invariably refer to the same specific country.

Finally, one also cannot assume that if two different codes are used in an EHR then they refer to different entities. Possible reasons for this are manifold. It may be that the most specific or detailed code is not always used when the same entity is referred to on successive occasions. A *colon polyp*, for example, when re-examined in the course of a follow-up visit where no change has been observed, might simply be referred to as *intestinal polyp*, or just *polyp*, and thus associated on successive occasions with different codes, even though the physician was fully aware that it was the same instance (the same polyp) that was being referred to.

It might also be that the polyp has become malignant, and then it will be assigned the code for *malignant neoplasm of colon*. Clearly, the relevant entity, i.e. the polyp, underwent changes. But it is still the same entity: its identity did not change. (In a similar way, persons undergo changes, grow older, lose hair, and so forth, but still remain the same continuant entities, preserving their identities over time.)

This preservation of identity in the presence of phenotypic difference is important in matters of prevention. As an example, there was a time when it was not yet common knowledge that particular polyps may deteriorate and become cancerous. Different surgeons may indeed have observed this process in particular patients; but they were not aware that by reporting what they had observed they would be contributing valuable scientific information. Suppose, now, that a statistical study targets patients suffering from intestinal polyps, the patients being selected on the basis of the presence of the code for *intestinal polyp* in their EHRs at some time t1. From the records of these patients one can also extract all disorder codes that are registered at times t later than t1. Imagine, however, that in a statistically significant portion of the latter, the code for *malignant neoplasm of colon* has been registered. As a result, taken over all records, one may conclude that the presence of an intestinal polyp is a risk factor for the appearance of a malignant tumor at some later time. But if, as is the case under current EHR regimes, one had no systematic means of recording that the very same (i.e. numerically identical) polyp—rather than a second polyp discovered elsewhere in

the colon—had turned malignant over time, then one would also not be led to postulate the removal of the initial polyp as the best prevention for malignancy.

A third reason why different general codes from a coding system may not automatically be taken to refer to different particular instances turns on the fact that a code may not suffice to describe a given instance appropriately. If, for example, one wants to use SNOMED-CT (v0301) to code a *closed pedicular fracture of the fifth cervical vertebra* then a single code is not available; to give a faithful description one must use instead the codes for both '*fracture of pedicle of cervical vertebra*' and '*closed fracture of fifth cervical vertebra*'. If, however, these codes are not entered in the EHR in such a way that it is clear that they refer to the same particular entity, then their presence might be taken incorrectly to refer to two different fractures.

Similar mistakes may arise also where the same class is represented twice by means of two different codes in a single coding system (e.g. in SNOMED-CT v0301, 41191003: '*open fracture of head of femur*' and 208539002: '*open fracture head, femur*') [2].

As an intermediate conclusion we can therefore state that, even where coding systems and terminologies provide rich vocabularies to describe the entities that exist in reality using general terms, they have as yet been associated with no mechanism to express *what those descriptions are about*, i.e. what entities in reality they refer to. This is not, be it noted, a criticism of existing coding systems. The latter were not, after all, designed to have such a mechanism in place. But it is our claim that such a mechanism is indispensable at the interface where coding systems meet the clinical record if we are to gain maximal advantage from coding efforts and from so-called formal representations and descriptions in EHR systems.

2. Referent tracking

Drawing upon our experience in EHR research and standardisation [3–5], and also from our philosophical research on universals and particulars [6], we introduced *referent tracking* as a paradigm under which it will become possible to refer explicitly to all of the concrete individual entities relevant to the accurate description of each patient's condition, therapies, and outcomes through the assignment of unique identifiers [7]. Such an identifier is called a *IUI*, for *Instance Unique Identifier*. This means that not only does the patient receive a IUI, but so also does the particular fracture he is suffering from, the particular bone that is fractured, and even, if the clinician finds this important, the particular pain the patient is experiencing in a certain time period or the particular document in which the pain is first recorded.

As such, referent tracking goes much further than current practices, under which entities are uniquely identified only when they belong to a restricted range of entity-types, including human beings (the patient himself, the physicians involved), buildings (the hospital in which the patient is treated), certain instruments and devices, and so forth. Moreover, where the majority of entities uniquely identified under current schemes are outside the patient (physicians, instruments, wards, and X-ray images), referent tracking extends the facility of unique identification also to the patient's body parts, the specific diseases he has suffered from, the symptoms he has exhibited, and so forth. It goes beyond established approaches also in the degree to which it takes seriously the notion of 'uniqueness.' For where, in many EHR systems, patients and physicians are uniquely identified only relative to some local context (for example of the hospital in which a given EHR-system is used), referent tracking aims for *global* uniqueness.

Note that IUIs refer to the real entities themselves out there in reality, and not to data about these entities. IUIs are the means whereby those constellations of particular entities (tokens, instances) in reality that are relevant to clinical care can be represented in an EHR in the same direct way in which the corresponding classes (types, universals) are already represented by means of clinical coding systems.

Thus IUIs are also *not the entities themselves*. This might seem obvious, but use-mention confusions ('Swimming is healthy and contains eight letters')—in which an entity in reality and its digital representation are confounded together—are abundantly present in the literature on knowledge representation in general and on concept-based terminology systems in particular [8]. (See the discussion of HL7 below.)

In the context of this paper we will use expressions of the form 'IUI-*uvwx*...' , where $u, v, \ldots$, will be substituted by numerical digits, to denote the identifiers themselves. We will then be able to discuss, for example, the length of a IUI, or the font in which it is written. Expressions of the form '#*X*,' in contrast, where '*X*' stands proxy for a IUI, will denote the corresponding individual entity in the real world. This will allow us to write for instance that IUI-3006 *refers* to the particular patient John, while #IUI-3006 *is* that particular patient.

The referent tracking paradigm distinguishes between *IUI assignment*, which is possible only in relation to entities that exist or have existed in the past, and *IUI reservation*, which is a provision made for entities, such as an X-ray ordered for tomorrow, that are expected to come into existence in the future. The order itself can have a IUI assigned already today, but for the resultant image one can only reserve a IUI at the time of ordering.

Note that IUI assignment or reservation does not by itself entail any assertion as to the class (or, since we take a position grounded in realism as a philosophical theory, the *universal* [9]) of which the particular in question is an instance. Thus we might assign a IUI to some syndrome of a given patient before we have any clear idea what sort of syndrome it is with which we are dealing. This facility, too, has no analogue in code-based EHR systems as currently constituted and that resort to invent new classes such as 'unknown syndrome.'

In this paper we explore ways in which the referent tracking paradigm can be implemented in the healthcare environment. Our hypothesis is that, once the right infrastructure is in place, the burden on clinicians and nurses (or on whomever is assigned the task of registering patient data) will be not significantly greater than under existing strategies for data entry—but that the benefits, in terms of semantic interoperability of computer systems and also in terms of patient management, cost containment, epidemiology and disease control, as well as for the advance of science in the domain of biomedicine, can be enormous.

3. Overall architecture of a referent tracking system

3.1. Dealing with referents

The purpose of a referent tracking system (RTS) is, as its name suggests, to keep track of *referents*. Referents are entities that exist in reality, i.e. in the spatiotemporal world that surrounds us. Most referents are *particulars*, examples being a copy of the journal in which this paper is published, its authors, those of its readers who find its ideas appealing, as well as those readers who doubt the existence of physical reality (and hence their own existence). Other referents are *universals*, examples being *journal*, *paper*, *person*, and so forth. According to the philosophy of realism, universals are as real as particulars. (Realism is thus opposed both to nominalism, which denies the existence of universals, and to conceptualism, which asserts that what is general in reality exists only in the minds of particular concept-using subjects.) Referent tracking deals primarily with the tracking of particulars (so that even nominalists and conceptualists may be able to take advantage of its potential). In this paper, we focus on the tracking of particular referents in the context of maintaining EHRs, but the paradigm is clearly applicable in other contexts as well. We also assume without further consideration that the identification of universals will be taken care of in adequate ontologies [10].

An RTS will contain information about particulars. The users who enter this information will be required to employ IUIs in order to assure explicit reference to the particulars about which they are providing information. Thus the information that is currently captured in the EHR by means of sentences such as: ''this patient has a left elbow fracture,'' would in the future be conveyed by means of descriptions such as ''#IUI-5089 is located in #IUI-7120,'' together with associated information to the effect that ''IUI-7120'' refers to the patient under scrutiny, and ''IUI-5089'' to a particular fracture in patient #IUI-7120 (and not to some similar left elbow fracture from which he suffered earlier). The RTS must correspondingly contain information relating particulars to universals, such as ''IUI-5089 **instance_of** fracture'' (where 'fracture' might be replaced by a unique identifier pointing to the representation of the universal *fracture* in an ontology) [6]. Of course, EHR systems that endorse the referent tracking paradigm should have mechanisms to capture such information in an easy and intuitive way, including mechanisms to translate generic statements into the intended concrete form, a form which may itself be operative primarily behind the scenes, so that the IUIs themselves remain invisible to the human user. One could indeed imagine that natural language processing software will one day be in a position to replace in a reliable fashion the generic terms in a sentence ('John's mother,' 'John's pacemaker') with corresponding IUIs for the particulars thereby denoted, with manual support in flagged problematic cases. This corresponds on the level of particulars to what users already expect from EHR systems on the level of universals in supporting entry of codes or terms from coding systems.

3.2. Minimal requirements for referent tracking

At least the following requirements have to be addressed if the paradigm of referent tracking is to be given concrete form:

- a mechanism for generating IUIs that are guaranteed to be unique strings;
- a procedure for deciding what particulars should receive IUIs;
- protocols for determining whether or not a particular has already been assigned a IUI (except for some exceptional configurations that are beyond the scope of this paper, each particular should receive maximally one IUI);
- practices governing the use of IUIs in the EHR and in clinical documentation and research in general (issues concerning the syntax and semantics of statements containing IUIs);
- methods for determining the truth values of propositions that are expressed through descriptions in which IUIs are used;
- methods for correcting errors in the assignment of IUIs, and for investigating the results of assigning alternative IUIs to problematic cases;
- methods for taking account of changes in the reality to which IUIs get assigned, for example when particulars merge or split.

An RTS can be set up in isolation, for instance within a single general practitioner's surgery or within the context of a hospital. The referent tracking paradigm will, however, serve its purpose optimally when it is used in a distributed, collaborative environment. One and the same patient is often cared for by a variety of healthcare providers, many of them working in different settings, and each of these settings may use its own separate information system. These systems contain different data, but these data often provide information about the same particulars. Under the current state of affairs, it is very hard, if not impossible, to query these data in such a way that, for a given particular, all information available can be retrieved. With the right sort

of distributed RTS, such retrieval becomes in very many cases a trivial matter.

3.3. Services to be provided

The system we have in mind should offer at least three services: to generate unique identifiers to be used as IUIs, to keep track of the IUIs generated and to provide access to the IUIs stored.

As to the first, we shall see below that methods already exist for generating globally unique random alphanumeric IDs of the sort required. Thus this aspect of the referent tracking system poses no new technical or theoretical problems.

The second, which involves what we shall refer to as the *IUI-repository*, is the most crucial service to be provided by an RTS, and consists in keeping track of the identifiers assigned to already existing entities, or reserved for entities that are expected to come into existence in the future. It will do this in such a way that each IUI represents exactly one particular, and that no particular is referred to by more than one IUI. These two requirements are not easy to fulfil, since both depend on the ability and willingness of users to provide accurate information. This, however, introduces no problems different in principle from those already faced by the users of existing systems when called upon to provide information of a non-trivial and occasionally sensitive sort.

The third service, here called the *referent-tracking database* (RTDB), should provide access to the information entered into given EHRs about the particulars referred to in the IUI-repository. The IUI repository is an inventory of concrete entities that have been acknowledged to exist, and, consequently, of what IDs to use if one wants to refer to them. The RTDB, in contrast, is an inventory of descriptions concerning the features and interrelations of these entities and the ways they change in the course of time. The RTDB, too, does not need to be set up as a single central database (in which case it would be some sort of data warehouse); rather, it could rely on any paradigm for distributed storage.

The primary role of the RTDB is to keep track of the features of given particulars and of their relationships to other particulars as they change through time, and of the assertions that have been made about such particulars, including those assertions which have been shown to be false (which are stored for the medico-legal purposes of providing an audit trail). It has an important role also in helping users to determine whether particulars they encounter for the first time have been registered already in the IUI-repository, or whether a new IUI must be created for use in new descriptions. To be sure, this places some additional burden on the person who has to enter information; but time perceived as being lost at this stage will be recovered when searching for information thereafter. Moreover, the additional burden will be incurred almost always in relation to those cases where the particular in question enjoys a high degree of salience in the medical history of the patient in question.

4. The generation of IUIs

Several schemes for generating strings that are guaranteed to be unique are already in use, the most prominent of which is Microsoft's Globally Unique Identifier paradigm (GUID), which implements UUIDs (Universally Unique IDs) as defined by the Open Software Foundation in the specification of its Distributed Computing Environment [11]. The advantage of the GUID paradigm is that unique identifiers can be generated easily on any machine with a network card, without the need to resort to a central authority to guarantee uniqueness.

UUIDs have recently been standardized through ISO/IEC 9834-8:2004, which specifies format and generation rules that enable users to produce every 100 ns 128-bit identifiers which are either guaranteed to be, or have a high probability of being, globally unique [12]. The standard also specifies the procedures for the operation of a Web-based Registration Authority for UUIDs. Although some older versions of UUID generating algorithms may produce IDs that contain meaningful information (such as the MAC address of the machine used to generate the ID), recent versions no longer exhibit this behaviour.

UUIDs have hitherto been used only for unique identification of software components such as the pop-up windows generated in the course of a program's execution. But there is no reason why they should not be used also to identify particulars in the world outside the machine. In the specific case of health-related particulars, ethical, safety and security considerations might require *certification* of their uniqueness, given that there is still the risk of double generation under the 100 ns time limit in case UUIDs are used, or in case other schemes would be preferred such as for example the Object Identifier standard (OID) [13]. To that end, the medical community may want to install an authority that not only registers IUIs, but also certifies the uniqueness of the strings to be used within a given IUI-repository and guarantees that the assignments claimed to have been made by given authors were indeed made by those authors. This can be compared to the services offered by trusted third parties in private key management for asymmetrical encryption purposes [14]. Moreover, central registration in some form will in any case be necessary if we are to fulfil a requirement (explained in the next section) to the effect that no particular be assigned more than one IUI.

Note that an IUI does not carry any information as to which particular it refers to. It is a "meaningless" alphanumeric string. As explained above, information about what IUIs actually stand for is to be found in the RTDB.

This scenario has a number of advantages over the proposals, sometimes advanced on behalf of existing EHR-systems, of creating 'meaningful' identifiers by combining existing patient-IDs with modifiers such as '*right* [ ]' or '[ ] *of* [ ]' (as in '*right* leg' or 'kidney *of* patient John') to enable

unique reference to corresponding particulars on the side of the patient. While such proposals might seem at first glance to be easy to realise, they have in fact nowhere been implemented in a systematic way. They face problems, first of all, because they are subject to familiar scope confusions introduced with the use of general terms (does 'leg' refer to the lower leg or to the entire lower limb structure?, does 'third bed from left' refer to the bed or to its current occupant?). They can also lead to inconsistencies, for example when patient A's left kidney is replaced by that of patient B after a kidney transplantation. Problems arise, too, because for each given particular there will standardly be a plurality of different definite descriptions, all of which identify the particular uniquely but via modification of different general terms. This creates an obstacle to reasoning about the corresponding particulars because the information supplied in each separate description will in general not be derivable from that supplied in the others (compare 'morning star,' 'evening star'). Definite descriptions are affected also by problems referred to above, turning on the fact that different hospital institutions may have different coding habits when it comes to using general terms, or on the fact that a plurality of general codes may be needed to specify an instance in an adequate fashion in a given terminology. We conclude that, since the information carried by composed identifiers would in any case be contained in expressions in the RTDB, the referent tracking paradigm is to be preferred to any proposed solution based on descriptions formed via modification of general terms.

5. To assign or not to assign

As already noted, IUIs should be *assigned* exclusively to entities that exist or have existed in the past, as contrasted with those cases where entities are expected to come into existence in the future, where IUIs can only be *reserved*. This formal difference is an intrinsic part of the referent tracking paradigm. Although each particular is, by definition, a unique entity, particulars are standardly not such as to have uniquely identifying labels already attached. Newborns in the USA are assigned Social Security numbers only as a result of an application procedure. IUI assignment, in general, is typically an act carried out by the first cognitive agent who feels the need to acknowledge the existence of a particular it has observed (or has information about of a sort that is analogous in its reliability to that which is gained through observation). In the healthcare environment the assigning entity will typically be a person (clinician, nurse, patient); but it might also be a device, as for example when radiographic films are manufactured in such a way that each film is tagged with its own IUI automatically; analysis software that operates on digital images might automatically assign IUIs to specific configurations found therein, such as fracture lines or coin lesions.

From a logical perspective, each act of IUI assignment rests on a complex belief (a presupposition), on the part of the cognitive agent involved, to the effect that the following three propositions are true:

1. this particular (here before me now) exists;
2. this particular has not yet been the object of IUI assignment;
3. the string that functions as IUI for this particular has not been used thus far for any other entity.

The agent in question acquires this complex belief on the basis of prior consideration of the particular in question and his knowledge of the workings of the pertinent referent tracking system in light of its fulfilment of four criteria—of existence, uniqueness, no prior assignment, and salience—set out in the section which follows.

5.1. Criteria for IUI assignment

5.1.1. Existence

The first criterion which needs to be fulfilled before a IUI can be assigned is: "*Does what I want to assign a IUI to exist?*" Only if the answer to this question is 'yes' is assignment allowed. Where the particulars in question are a patient's body parts or objective signs such as skin lesions, checking existence is trivial. Other cases—such as a patient's symptoms as subjectively reported—will be more problematic. Certainly when a patient complains about a headache, then his making this complaint is a particular utterance (an event) to which a IUI can unproblematically be assigned. But that event is of course distinct from the particular which is the headache itself.

As stated already above, consideration of the exact nature or type of a particular (of the universals it instantiates) does not play a role at the stage of checking its existence. That we do not know all that there is to know about a given particular is an epistemological issue, which has no consequences for the ontological status of the particular itself. (If we are realists about the future, so that we view the existence of objects in the future as analogous to that of objects in the past and present, then the difference between assigning and reserving an IUI may be a special case of this epistemological issue.) Suppose that patient #IUI-001 has been stung by a particular bee from a particular swarm (#IUI-2345) consisting of bees of a sort to which he is allergic. Our not knowing precisely which bee is responsible for the sting does not prevent us from assigning a IUI, e.g., IUI-567, to this bee. We can then state that #IUI-567 is a member of #IUI-2345 and that it stung #IUI-001 at time t1. The same strategy may be used to refer to the particular pain attack (from a series), or to the particular pustule (from a group of severely inflamed acne pustules on the patient's face), that finally made the patient decide to consult a physician.

5.1.2. Uniqueness

The second criterion which needs to be fulfilled requires that it be established that the particular to which one wants to assign a IUI is not the same as some particular whose

existence has already been acknowledged (whether or not it has already received a IUI). The classical example used in the philosophy of language is the planet Venus [15]. For a long time, people believed in the existence of two distinct heavenly bodies, one visible in the firmament in the morning, the other in the evening. In reality, however, it was the very same planet that was being perceived on both sets of occasions. The same situation may be encountered in healthcare, for instance when, in the course of a patient's medical history, a disease is assumed to be the cause of certain manifestations which then disappear, a second disease is a few years later assumed to be the cause of certain other manifestations which also disappear, until it is finally established that one and the same (i.e. numerically identical) underlying disease, for example multiple sclerosis, had been causing all observed manifestations from the very beginning. Another example is that of several different patients all of whom have been bitten, successively, by the same dog. Of course, all these cases pose certain epistemological problems: an observer cannot be sure about whether what he observed is actually true; he might get things wrong. Hence the referent tracking paradigm has to include a facility for dealing with mistakes. (This will not be dealt with further here, but some initial remarks can be found in [10].)

5.1.3. No prior assignment

The third criterion concerns whether or not the particular whose existence and uniqueness has been determined has already been assigned a IUI, since—drawing on familiar arguments which were used to justify the introduction of unique patient identifiers [16]—our paradigm insists on at most one IUI per particular. This is because information in different (or even in the same) EHRs might otherwise not be interpretable as pertaining to the same entity. This is not to say that other, temporary, IDs might not be produced for pragmatic reasons, for example when the referent tracking system itself is for whatever reason off-line, or when data has to be entered under extremely urgent circumstances and one does not have the time to perform an adequate search in order to establish in reliable fashion that all criteria for IUI assignment have been met. The resultant IDs are not full-fledged IUIs, however, and steps should be taken to ensure that they can always be replaced by IUIs resulting from proper assignment procedures at a later stage.

In order for a cognitive agent to assess whether or not a particular entity did already receive a IUI, the agent must read and understand the pertinent descriptive information in the RTDB. While human agents will have to rely on textual descriptions, or on translations into human readable form of descriptions offered in a pre-specified formal syntax [17], software agents could use the formal representations directly. Imagine a physician performing a follow-up colonoscopy on a patient, who thereby sees a polyp. The physician will in many cases already know that the RTDB contains an entry for this polyp in an earlier stage of its existence because he himself, or a colleague with whom he communicated, was responsible for entering the corresponding information. In other cases, however, he must use available evidence to decide whether or not the polyp that he is looking at is indeed the very same as one previously seen. Thus he must examine associated data that describe size, location, appearance, etc. of the polyp—and hence the importance of incorporating IUIs referring to related entities, in formulating descriptions. Clearly, in establishing whether or not he is reporting on something new, the physician will in some cases be confronted with considerable epistemological difficulties. We believe, however, that in an era of increasingly personalized medicine, these are difficulties which must in any case be confronted by any EHR regime adequate to the information-processing needs of the future, and thus they are independent of the paradigm of referent tracking as such.

5.1.4. Salience

The fourth and final criterion concerns whether a given particular is, from the point of view of clinical care, sufficiently salient to justify the assignment of its own IUI. Here we can point to the clear distinction drawn already in everyday life between those entities that are considered sufficiently important to receive their own unique IDs (usually proper names rather than numbers)—for instance children, pets, large farm animals, yachts, valuable diamonds—and those considered not so important (snowstorms, molecules, ocean waves). In everyday life it is usually *continuants*, i.e. entities that preserve their identity through time (and thus are wholly present at every time during the course of their existence, even though they may gain or lose parts from one time to the next), that are given proper names. In healthcare, however, it may be that particular *occurrent* processes are of equal importance. Thus it may be crucial for medico-legal and other purposes that particular acts of giving injections, removing or transplanting organs, reducing fractures, and so forth are uniquely identified. For it might well be that a specific act performed at time t1 leads to consequences different from those yielded by a second, exactly similar act at time t2, or that the two lead to the same immediate consequences but with different long-term effects.

Bearing in mind the progressive advantages (for example to statistical reasoning and as a prophylactic against litigation) of ever larger repositories of clinically relevant data, and also the ever decreasing costs of computer memory, we propose accordingly that all entities to which one would standardly refer, either individually or generically, under the current rules and practices of clinical record-keeping, should receive Instance Unique Identifiers. Because such rules and practices differ from one institution or physician to the next, IUI assignment will not everywhere be implemented in the same way or to the same extent. Thus for EHR systems built primarily around notes in natural language, which offer only minimal facilities for structured reporting, we would expect

a less demanding policy. However, we advocate the principle that *if* a particular has already been assigned a IUI by somebody else, then it should be used also in these less advanced systems.

5.2. Publishing IUI assignments

When, after due consideration, a particular has been identified as requiring a IUI, then a thus far unused alphanumeric string will be generated by the ID generator and an act of *assignment* will be carried out (analogous to an act of baptism). This creates a IUI out of the generated string, by attaching it to the particular in question [18]. Three factors can be distinguished as structural elements involved in such an assignment act:

1. the generation of the relevant alphanumeric string;
2. its attachment to the relevant object;
3. the publication of this attachment.

The resulting IUIs will, together with certain further types of associated information, constitute the *IUI-repository*. The units deposited in this repository can be represented as ordered tuples of the form

$$A_i = \langle IUI_p, IUI_a, t_{ap} \rangle$$

where IUI_p is the IUI of the particular in question, IUI_a is the IUI of the author of the assignment act, and t_{ap} is a time-stamp indicating when the assignment was made ('*A*,' here, stands for *assignment*.) In light of the need to resolve mistakes in IUI assignments, each such tuple will need to be complemented with meta-data recording by whom and at what time they were deposited in the system. These meta-data are ordered triples of the form

$$D_i = \langle IUI_d, A_i, t_d \rangle$$

where IUI_d is the IUI of the entity registering the IUI in the system, A_i is the information-unit in question, and t_d is a reference to the time the registration was carried out (which is also the time from which IUI_p can be used in descriptions concerning the particular in question). ('*D*,' here, stands for *deposit*.)

While neither t_{ap} nor t_d should be assumed to carry information about when the particular referred to by IUI_p started to exist nor about its continued existence, it can be inferred that IUI_p did not start to exist at a time later than t_{ap}.

5.3. Management of assignments

Both the A_i and D_i should be stored in the IUI-repository in such a way that they can be accessed by software applications. While the repository might be a single centrally maintained database, since its contents are data, they may also be addressed by appeal to the LSID paradigm [19]. Although 'LSID' is an abbreviation for *Life Science Identifiers*, the LSIDs are in fact more properly conceived not as identifiers but as addresses: they inform a software program where it can find data.

It should be a requirement for all systems that are part of a referent tracking environment that the *A*-(assignment) tuples should be registered in the IUI-repository in association with corresponding *D*-(deposit) tuples. This is a necessary (though not a sufficient) condition for ensuring that no particular is given two distinct IUIs. Ideally, the A_i-tuple should be entered into the system as soon as an assignment act has taken place, and the corresponding D_i-tuple as soon as possible thereafter (normally just a short time later). Clearly, measures should be implemented to prevent content being deliberately entered that is based on false information. We suggest that any computer system contributing to the referent tracking architecture should require each user to log in in such a way that his own IUI becomes associated automatically with all the data he enters. This brings also the advantage that when an A_i-tuple is entered whose IUI_a component is the IUI of the user, then the corresponding D_i-tuple will be generated automatically.

Additional rules for IUI assignment may be implemented also. Thus one might require that the patient (or several patients, for particulars of certain kinds) authorises those who are allowed to make IUI assignments for particulars that concern him, who may or may not correspond to the persons authorised to manage his EHR. This '*need to enter*' policy is similar to the '*need to know*' policy installed in those institutions where authorisation to use a hospital information system does not entail authorisation to access *all* its information. This requirement would also limit the (from the point of view of the patient) uncontrolled accumulation of information about his health. Another rule might state that when clinicians do not enter patient data directly into the system but use some form of transcription, then they may authorise specific transcribers to register IUI assignments on their behalf, and so forth.

Importantly, the RTS should not allow the deletion of already existing content: neither the existence of a particular, nor the act of IUI assignment (itself a particular in its own right), can be undone. Of course, errors may occur and one might discover that a putative particular did not exist, in which case the error must be corrected. But when errors in IUI-assignment are corrected, then the original assignment information must be preserved, albeit in a form which ensures that those who use the information are aware of its erroneous nature.

6. Using IUI's in EHR statements

Once a IUI is registered in the referent tracking environment by means of *A*- and *D*-tuples, it can be used in descriptions of relevant facts or hypotheses about a patient's medical condition, his treatment, risk factors, and so forth. Descriptions may be directly about the particular itself, but also about other particulars that stand in some relation to it. Thus, if IUI-924 refers to patient #IUI-

0067's temperature, then it may be asserted that at time t1, #IUI-924 had the value 37.6 °C. Such a statement may be expanded to include also information about who performed the measurement (#IUI-3456), using what instrument (#IUI-4109), and so forth. The statement is then not directly about #IUI-3456 or #IUI-4109; but still it tells us indirectly as much about these particulars as it does about the patient's temperature at that specific time, and hence might provide useful management information. (For example a device that has been used a certain number of times might require maintenance; or statistically significant differences might be detected in the measurement data obtained using similar devices or by specific persons).

What should be included in descriptions of any given particular is not something that is dictated by the referent tracking paradigm itself. Users, when entering data, may follow the recommendations for good clinical registration issued by relevant bodies, openEHR archetypes being just one example [20]. As with statements about IUI assignments, and in line with recommendations pertaining to EHR standards such as ENV 13606 [21], associated descriptions should also as soon as possible after they are entered in the system be registered as having a precise author and tagged with data stating who made the information available and at what time.

The format in which information can be entered in a specific EHR system depends on the facilities the EHR system offers. Many systems do not allow a *formal* notation but expect data to be entered rather by means of natural language text. In the following paragraphs we first discuss some strategies for using the referent tracking paradigm in such text-based systems, before moving on to discuss ways in which data may be entered in more structured EHR-environments. For both types of system, we assume that the user is authorised to have access to the IUI-repository and is able to view the information about given particulars that is available through the RTDB.

6.1. IUIs in text-based EHR systems

The problem of how to deal with references to particulars in text-based EHR systems is not significantly different from the problem of how to deal in such systems with codes from concept-based terminologies such as SNOMED-CT or classifications such as ICD-9.

In the worst case, a system foresees nothing at all in the way of coded data entry. Typically, such systems allow you to type in free text in fields labelled *complaints*, *symptoms*, *diagnosis*, and so forth. All the user can do in order to enter codes—and IUIs under our paradigm—is to type them into the same data entry field as he types the free text. For example, he might just write them at the end of each natural language statement, using some syntactic convention to separate them from the text itself. A suitable interface would then need to be able to inform the RTDB that the sentence in question contains information about the particulars referred to by the IUIs listed.

For example in

Open left elbow fracture was reduced with pins in 1984 (IUI-5089, IUI-1002, IUI-4900)

IUI-5089 might refer to the fracture, IUI-1002 to the reduction and IUI-4900 to the pins, though this fact cannot be derived from the statement itself.

Theoretically, it would be possible to use a more elaborate syntax, such as that used in Cassandra-tagging along the lines proposed in [22]. The mentioned example would then be written (using indentation for better display) as

```
(   (open elbow fracture)IUI-5089
   {(reduced)IUI-1002
      {[with] (pins)IUI-4900}
      {[in] (1984)}
   }
)
```

Note that the individual phrases in the example above—'*open elbow fracture*,' '*reduced*,' etc.—do not give a uniquely identifying description of the particulars that are referred to by the IUIs that follow these phrases syntactically. For #IUI-5089 this is obvious: no elbow fracture is "just" an elbow fracture. It must be of some specific elbow, namely either the left or the right elbow of some particular patient (the patient whose case is being described). But it might be that #IUI-5089 is already described elsewhere in more detail (detail that might be found by searching the RTDB), or that this detail is not known (for example if the registration is entered in the course of a patient anamnesis and the patient does not remember whether the fracture was on the left or on the right). Even if the phrase would have read '*open left elbow fracture*,' then it still would not uniquely identify the fracture. An identifying description of the fracture would be obtainable for instance via cross-reference to IUI-1002 as well, and this only if #IUI-1002 is that very precise reduction event in which #IUI-5089 figured as exclusive participant as the fracture being reduced.

The same example as above, but using SNOMED-CT codes instead of IUIs, might look like this:

```
(   (open elbow fracture)302232001
   {(reduced)122469009
      {[with] (pins)77444004}
      {[in] (1984)}
   }
)
```

where *302232001* is the SNOMED-CT concept code for the concept with the fully specified name '*elbow fracture—open (disorder)*,' *122469009* the code for '*reduction procedure (procedure)*' and *77444004* the code for '*bone pin, device (physical object)*'.

Of course, one cannot expect a clinician or nurse to enter patient data by means of a Cassandra-like syntax. Renderings of the types proposed should rather be the outcome

of subjecting free-text statements to automatic natural language analysis [23,24]. An intermediate solution would be to allow the text editor used for entering natural language sentences in the EHR to incorporate hyperlinks through which the user could enter IUIs associated with the linked phrases. With a user-interface of this type, it would thereafter be possible to right-click on the hyperlinked phrases in order to launch a query to the RTDB that would then return all or targeted portions of the information related to the particular under scrutiny.

6.2. IUIs in EHRs incorporating formally structured statements

EHR systems incorporating record architectures, such as those proposed by GEHR [25], OpenEHR [26] or CEN ENV 13606 [21], would be almost ideally suited to the referent tracking paradigm. Although none of these architectures currently take particulars properly into account (because they are all built on the basis of the concept paradigm), the modifications that would need to be made are minor. In the case of CEN ENV 13606, for example, it would require an additional *compound data type* to be defined in order to make it formally clear that the content of a particular *data item*, i.e. one of the architectural components defined by the given standard, is a IUI, rather than a proposition. The modification we propose would then allow instance data to be exchanged between EHR systems that endorse the CEN ENV 13606 standard in such a way that formal reasoning about these data, including reasoning applied to data drawn from different systems, would become more reliable.

Particularly interesting from the point of view of the information they provide are descriptions in EHRs stating who or what a particular entity under scrutiny actually *is*, for it is these which are most relevant for determining whether a particular for which IUI assignment is being considered has been already registered in the IUI-repository. The strongest statements would be those that would enable an interpreter to point to the particular in question without the possibility of error. If all particulars would carry their IUIs with them (as it were indelibly attached), then the (universally applicable) statement "*#IUI-xyz is that particular which carries IUI-xyz*" would be all that would be needed to identify particulars unambiguously. While this facility is unfortunately available in the healthcare environment only in rare cases (for instance prosthetic devices with unique manufacturer-assigned serial numbers), it is beginning to be exploited in systematic ways in other domains. Thus it is reflected in hard- and software implementations exploiting RFID (Radio Frequency Identification), which yields statements of the form: "*#IUI-xyz is the particular that produces IUI-xyz when probed by an appropriate sensor*" [27,28].

While statements of the latter sort provide identity criteria for particulars, they are not informative with respect to which universals the particulars instantiate, i.e. to what *kinds* of particulars they are. For this, an ontology is required, which means a representation of whatever is the pertinent domain of reality which

(1) reflects the universals instantiated by the particulars and the relations between both universals and particulars in that domain in such a way that there obtains a systematic correlation between reality and the representation itself,
(2) is intelligible to a domain expert, and
(3) is formalised in a way that allows it to support automatic information processing.

The coding systems in common use, which we here refer to by means of the term "*concept-based systems*," do not meet these requirements, since they provide no reliable means to reason from information about concepts to information about the corresponding particulars in reality. Systems that do conform to this definition are BFO [29], the OBO Relation Ontology (RO) [6], and the Foundational Model of Anatomy (FMA) [30]. Ontologies of this kind contain relationships between universals that are formulated in such a way that they can be used in automatic fashion to describe also relationships between the corresponding particulars [31].

6.2.1. Relationships between particulars

RO distinguishes (1) relations that obtain between particulars, (2) relations that obtain between particulars and universals, and (3) relations that obtain between universals. In this paper, we will use **bold** type to indicate relations of types (1) and (2), and *italics* to pick out relations of type (3). Type (1) relations can be used to formalise descriptions in an EHR system which assert relationships between precisely those particulars that are relevant to the given patient, rather than between the corresponding universals, and so enables such descriptions to be much more narrowly targeted: #IUI-1921 (the first author's left testicle) is the left testicle not just of some instance of *human being*, but of the very precise particular #IUI-0945.

RO does justice also to the fact that relationships such as parthood have distinct properties at the particular and at the universal levels [32]. Thus from the statement "#IUI-1921 **part_of** #IUI-0945 at time $t1$," one can infer that "#IUI-0945 **has_part** #IUI-1921 at time $t1$," while a similar conclusion on the universal level does not hold. Thus "left testicle *part_of* human being" does not entail that "human being *has_part* left testicle" under the usual interpretation given to such a proposition, namely that for all human beings h there exists some left testicle t such that h **has_part** t—because there are humans who do not have a left testicle, most of them being female. (We here leave aside the fact that non-human mammals also have testicles.)

Note that, in representing the relations that obtain between continuants such as #IUI-0945 and #IUI-1921, the factor of time may not be left out of account. It might indeed be that at a time later than $t1$, the **has_part** relation

between #IUI-0945 and #IUI-1921 no longer holds. Such considerations, too, do not arise at the level of relations between universals, and they are not taken account of in established concept-based systems.

6.2.2. Formal representation of relationships between particulars

Descriptions which express relationships amongst particulars we will refer to as *PtoP*—**p**articular to **p**articular—descriptions. Here again we can distinguish a number of structural elements:

1. an authorized user observes one or more objects which have already been assigned IUIs in the referent tracking system (RTS) in hand,
2. the user recognizes or apprehends that these objects stand in a certain relation, which is represented in some ontology o,
3. the user asserts that this relation obtains and publishes this assertion by entering corresponding data into the RTDB.

This relationship (R-) data will then take the form of an ordered sextuple

$$R_{\mathrm{i}} = \langle IUI_{\mathrm{a}}, t_{\mathrm{a}}, r, o, P, t_r \rangle$$

where

- IUI_{a} is the IUI of the author asserting that the relationship referred to by r holds between the particulars referred to by the IUIs listed in P,
- t_{a} is a time-stamp indicating when the assertion was made,
- r is the designation in o of the relationship obtaining between the particulars referred to in P,
- o is the ID of the ontology from which r is taken,
- P is an ordered list of IUIs referring to the particulars between which r obtains, and
- t_r is a time-stamp representing the time at which the relationship was observed to obtain.

P contains as many IUIs as are required by the arity of the relation r. In most cases, P will be an ordered pair which is such that r obtains between the particulars represented by its first and second IUIs when taken in this order. As with A-tuples, so each R-tuple must be accompanied by a corresponding D-tuple capturing when it was deposited in the referent tracking system.

From the example used earlier we could then derive the following tuples:

⟨*IUI-6231*, *18/04/2005*, **has-participant**, *RO*, ⟨*IUI-1002*, *IUI-5089*⟩, *1984*⟩
⟨*IUI-6231*, *18/04/2005*, **has-participant**, *RO*, ⟨*IUI-1002*, *IUI-4900*⟩, *1984*⟩

where #IUI-6231 is the author asserting the obtaining of the relationship **has-participant**, which is taken from RO. Note that RO currently has only two relationships to describe ways entities may participate in events, and thus it needs to be extended, for example in order to make it possible to distinguish between the role played by #IUI-5089 (the fracture mentioned before) which is that of *theme* and the role played by #IUI-4900 (the pins used) which is that of *instrument*, in treatment event #IUI-1002 [33].

As in the case of the A- and D-tuples introduced above, so also the R-tuples are presented using an abstract tuple syntax. It is not however anticipated that such tuples should be entered in this form directly by end-users. Rather appropriate user-interfaces would take care of corresponding technical transformations behind the scenes.

6.2.3. Relationships between particulars and universals

The second type of information that can be provided about a particular concerns what universal within an ontology it instantiates. Here, too, time is relevant, since a particular, through development, growth or other changes, may cease to instantiate one universal and start to instantiate another: thus #IUI-0945 changed from *foetus* to *newborn*, and from *child* to *adult*.

Descriptions of this type (which we will refer to as *PtoU* entries—for: **p**articular to **u**niversal) can be represented by ordered tuples of the form

$$U_{\mathrm{i}} = \langle IUI_{\mathrm{a}}, t_{\mathrm{a}}, \mathbf{inst}, o, IUI_{\mathrm{p}}, u, t_r \rangle$$

where

- IUI_{a} is the IUI of the author asserting that IUI_{p} *inst* u,
- t_{a} is a time-stamp indicating when the assertion was made,
- **inst** is the designation in o of the relationship of instantiation,
- o is the ID of the ontology from which *inst* and u are taken,
- IUI_{p} is the IUI referring to the particular whose *inst* relationship with u is asserted,
- u is the designation of the universal in o with which IUI_{p} enjoys the *inst* relationship,

and

- t_r is a time-stamp representing the time at which the relationship was observed to obtain.

Note that it is necessary to specify from which ontology *inst* and u are taken (and precisely which *inst* relationship in those cases where an ontology contains several variants [34]). Such specifications will not only ensure that the corresponding definitions can be accessed automatically, but also facilitate reasoning across ontologies that are interoperable with the ontology specified.

7. IUIs in relation to concept-based systems

In Section 6.2 we required the designations of relationships and universals used in statements describing proper-

ties of particulars to be taken from ontologies rather than from concept-based systems. There are many reasons for this, including:

- the relationships between particulars and the corresponding concepts are often left obscure in such systems [8],
- the relationships themselves are inadequately defined,
- there is an inconsistent reading of statements with respect to existential or universal quantification [35],
- ontology and epistemology are mixed together in inappropriate ways [36].

We do not blame the authors of such systems for these inadequacies. Everything man-made must be expected to contain mistakes; and realist ontologies, too, will not be error-free. What we question, rather, is the unprincipled way in which such systems have been put together [37,38]. The good news, on the other hand, is that, as we will explain below, when they are used in conjunction with an RTS some of these systems may be transformed into sound ontologies of the sort which will be free of at least many of the given types of errors.

7.1. How concept-based systems can help in referent tracking

Even in their current form concept-based systems can play a useful role in the context of the referent tracking paradigm, since the latter already brings many of the advantages to be gained through the enforcement of sound ontological principles.

Tuples similar in form to *PtoU* tuples, but in which *u*, i.e. the reference to a universal from an ontology, is replaced by a reference to a concept from a concept-based system, would be useful for searching the RTDB. Of course, the relationship to be used cannot be some variant of '**instance_of**' since the standard definitions in use for '*concept*' (such as '*unit of knowledge*' [39] or '*unit of thought*' [40]) disallow most particulars from being declared as instances of concepts. (Instances of concepts would be, perhaps, contents of a knowledge base under the first definition, or ideas in people's minds under the second.) But #IUI-1921 (the first author's left testicle) will never be an instance of a concept, however the latter notion might be defined or whatever might happen to that particular in the future. (Thus it will never itself be a unit of thought, though it might perfectly well be something towards which a unit of thought is intentionally directed.) Hence we prefer to bring concept-based systems back to the task for which they were originally designed, namely to assist specialists in a given domain in obtaining a better grasp of the variety of terms in use in that domain for purposes of communication in a particular language such as English or French. What we shall refer to as *PtoCO* tuples (**p**articular to **co**ncept code) will then have the form

$$Co_{\mathrm{i}} = \langle IUI_{\mathrm{a}}, t_{\mathrm{a}}, cbs, IUI_{\mathrm{p}}, co, t_r \rangle$$

where

- IUI_{a} is the IUI of the author asserting that terms associated to *co* may be used to describe *p*,
- t_{a} is a time-stamp indicating when the assertion was made,
- *cbs* is the ID of the concept-based system from which *co* is taken,
- IUI_{p} is the IUI referring to the particular which the author associates with *co*,
- *co* is the concept-code in the concept-system referred to by *cbs* which the author associates with IUI_{p}, and
- t_r is a time-stamp representing a time at which the author considers the association appropriate

Such tuples are to be interpreted as providing a facility equivalent to a simple index of terms in a work of scientific literature. The "annotation" of an entry in a database by means of a term from a controlled vocabulary such as the Gene Ontology [41] is a typical example. All that the information in such a tuple tells us is that, within the linguistic and scientific community in which the concept-system referred to by *cbs* is used, it is acceptable to use the terms associated with *co* to refer to the particular in question.

As an example, the tuple

⟨*IUI-0945*, *18/04/2005*, *SNOMED-CT v0301*, *IUI-1921*, *367720001*, *forever*⟩

tells us that the first author of this paper on April 18, 2005 asserted that his left testicle is, within the linguistic and scientific community in which SNOMED-CT is accepted for use, such that it may always be denoted by the phrase "*this left testis*," since the term "*left testis*" is recognised in SNOMED-CT v0301 as an adequate term for concept code *367720001*. (And also, however odd this may sound, "*this entire left testis*" is acceptable too, since "*entire left testis*" is an alternative term associated in SNOMED-CT with the same concept-code.) Furthermore, by taking advantage of the structure of (properly designed) concept-based systems, in which terms more generic than a given term (its ancestor terms in an *is_a* tree) are also acceptable, we can refer to #IUI-1921 in addition by using the phrases "*this testicle*," "*this male gonad*," "*this testis*," "*this genital structure*," ..., "*this physical anatomical entity*," and so on—though not (in spite of the fact that we would then still be progressing upwards in the SNOMED-CT *is_a* hierarchy) "*this SNOMED-CT concept*." This last *is_a* relationship is accordingly a mistake in the structure of SNOMED-CT.

For the reasons given already above, the relationships that are used in concept-based systems to associate concepts with each other are for the moment too imprecisely defined to be usable in describing relationships that obtain between corresponding particulars. They have some value, rather, in providing guidance on how to browse through the systems in question to find terms that can be used to de-

note related particulars in ways acceptable to given user communities.

7.2. How referent tracking can help concept-based systems maintenance

We believe that referent tracking, when properly used, can solve one of the most intractible problems in the domain of concept-based systems, namely how to map them amongst each other. Indeed, if referent tracking would be applied in a sufficiently large community, mappings between different terminologies would in course of time be generated as automatic by-products of the referent tracking effort. Systematic referent tracking would also solve the problem of how to reuse data that have been coded by means of older versions of specific systems, and also help in uncovering mistakes in such systems and in the application of such systems in given institutions, and so forth.

To see how this would work in more detail, imagine that patient #IUI-001 consults physician #IUI-201 working in hospital #IUI-211 in city #IUI-400, and that, in order to obtain a second opinion, the same patient thereafter consults #IUI-3900, a surgeon in the clinic #IUI-0098 in the same city. The EHR-system in #IUI-211 is not designed to work with formal ontologies, but it nonetheless has facilities to code data in detail using SNOMED-CT and to represent the data by means of *PtoCO* tuples. It is also connected to the same RTS to which the EHR-system of #IUI-0098 is connected. The latter, however, uses MEDCIN, a concept-based system of entirely different stamp [42]. Clearly, the entities described by both physicians will be the very same entities, so that many of their descriptions will contain the same IUIs; but different codes from two different concept systems will be used in their descriptions. If both physicians (or the coders or transcribers registering data on their behalf) have done their jobs properly, then their combined efforts will result in a mapping of a small portion of MEDCIN to a small portion of SNOMED-CT and vice versa. As more and more health facilities, some of them using the same, some using different, coding systems but servicing overlapping patient populations, become connected to the same RTS, there would in due course arise a massive pool of *PtoCO* tuples covering identical particulars. Such a pool of data could be mined automatically, not only with respect to the particulars that are described and about which relevant inferences can be made (as well as concerning the universals they are instances of), but also in order to uncover certain problems related to the concept systems employed in the creation of the data. These include problems connected with the various ways in which interdependencies between systems within a single organization are maintained (for example in the creation of mappings between a concept-based system designed for billing purposes and another designed for clinical coding), and also problems in single systems with respect to their adequacy to the task they are intended to perform. In fact, standard statistical techniques could be used to identify codes involved in specific patterns of misuse (or codes not used at all, and hence obsolete), and to identify clinicians who do not understand the intended meanings of certain codes (due to lack of training or because the concept system from which the codes are taken has poor documentation), as well as many other shortcomings in the process of documenting patient cases.

8. Discussion

8.1. Unique identification

The need for unique identification of clinically salient entities in a patient's documentation was recognized very early on in the history of medical informatics. The central idea in Weed's *Problem Oriented Medical Record* (POMR) is to organize all medical data around a problem list, thereby assigning each individual problem a unique ID [43]. Unfortunately Weed proposes to apply the IUI methodology only to problems, and not to the various particulars that cause the problems, are symptomatic for them, or are involved in their diagnosis or therapy. The same holds of Barrows and Johnson's problem-based approach, which suffers further from an ambiguity in its treatment of unique IDs, which sometimes seem to refer to problems themselves and sometimes to statements about such problems [44]. The argument often used in favor of a POMR is that it makes it possible to track a problem such as *chest pain* over time as it evolves into a problem of *angina*, from there into a problem of *myocardial infarction*, of *CABG* (Coronary Artery Bypass Graft), and so forth. However, we consider it wrong to use the labels '*chest pain*,' '*angina*,' '*myocardial infarction*,' and so on to denote what POMR defines as 'the problem.' Rather, these labels refer to very different kinds of particular entities that appear and disappear in the unfolding of the history of the problem, all of them related in various ways to another particular—namely the underlying disorder—by which the problem is caused. Hence we argue that an adequate POMR should embrace unique identifiers for particulars of all of these latter types too.

8.2. Unique identity in event-based approaches

Another example of an EHR regime involving the use of unique identifiers is that proposed by Huff et al. [45], who, refreshingly, take "*the real world to consist of objects (or entities)*." They continue by asserting: "*Objects interact with other objects and can be associated with other objects by relationships... When two or more objects interact in the real world, an 'event' is said to have occurred.*" Each event receives an explicit identifier, called an *event instance ID*, which is used to link it to other events (reflecting the goal of supporting temporal reasoning with patient data). This ID serves as an anchor for describing the event via a frame-representation, where the slots in the frame are name-value tuples such as *event-ID* = "#223," *event-family* = "diagnostic procedures," *procedure-type* = "chest X-

ray,'', etc. Via other unique IDs the framework incorporates also explicit reference to the patient, the physician and even to the radiographic film used in an X-ray image analysis event. Unfortunately, because they concentrate too narrowly on the events themselves [46], Huff and his associates do not allow explicit reference to *what is observed* during events. This is in spite of the fact that the very X-ray report that they analyse begins with the sentences: ''*PA view is compared to the previous examination dated 10-22-91. Surgical clips are again seen along the right mediastinum and right hilar region*.'' [45]. Because they have no means to refer directly to those clips, they must resort to a complex representation with nested and linked event frames in order to simulate such reference.

8.3. Unique identification in HL7

Unique identification is recognized as being of utmost importance also by HL7, and to that end Version 3 of HL7 introduced the *Instance Identifier (II) abstract data type* [47]. However, instead of having the corresponding strings conform to ISO/IEC 9834-8:2004 as under the regime proposed above, HL7 uses the ISO Object Identifier (OID) standard (ISO 8824:1990) as introduced originally in [13], which requires some authority for the assignment of each new identifier, the ID of the authority then being the first part of any OID it creates. The totality of all OIDs issued by that authority is called its '*branch*.'

The difference, here, is a matter of choice, the most important issue being the procedures installed to ensure that any ID (OID or IUI) refers to exactly one entity, and that any entity receives maximally one ID. The procedures selected by HL7 involve the establishment of an OID registry. HL7 then assigns OIDs in its own branch for HL7 users and vendors upon request as well as for public identifier-assigning authorities. HL7 requires registered OIDs to be used regardless of whether these organizations have other OIDs assigned from other sources [47]. Before assigning OIDs to entities of whatever sort, HL7 intends to investigate whether an OID has already been assigned by other sources. If this is the case, HL7 will record this OID in its catalog without assigning a duplicate one in the HL7 branch. HL7 intends to exercise diligence before assigning an OID in the HL7 branch to third parties, since given the lack of a global OID registry mechanism, one cannot be absolutely certain that there is no pre-existing OID assignment for such entities. Also, a duplicate assignment might happen in the future through another source. If such cases of duplicate assignment become known to HL7, HL7 will make efforts to resolve the situation. For continued interoperability in the meantime, the HL7-assigned OID shall be the preferred OID to be used by users of HL7. Clearly, analogous issues will have to be addressed under whichever approach to referent tracking is ultimately adopted.

What is disturbing in the HL7 approach is that it leaves unclear what exactly its OID machinery is supposed to identify.

The definition of *Entity* given in [[47]: HL7 RIM section 3.2.1] reads: 'a *physical thing, group of physical things or an organization capable of participating in Acts, while in a role*'; and we read in the accompanying *discussion* that: '*an entity is a physical object that has, had or will have existence*'.

In [47, Vocabulary], however, we learn from the explanation of the allowed value (*EntityDeterminerDetermined*) of the attribute *EntityDeterminer* that: '*The described determiner is used to indicate that the given Entity is taken as a general description of a kind of thing that can be taken in whole, in part, or in multiples*.' This contradiction—between *Entity* as thing and *Entity* as description of a kind of thing—exhibits the typical use-mention confusions outlined in [48]. Matters become even worse when we examine the other values the mentioned attribute can take, one being that of '*described quantified determiner*.' When the attribute has this value, we are told, this indicates '*that the given Entity is taken as a general description of a specific amount of a thing. For example, QUANTIFIED_KIND of syringe (quantity = 3,) stands for exactly three syringes*'. But a *particular entity* can hardly be taken to *be a general description* in any coherent ontology, and a *general description* itself can hardly be taken to satisfy HL7's definition of *Entity* as *a physical thing, or a group of physical things or an organisation*. Moreover, even if we ignore HL7's impenetrable use of language and its multiple levels of use-mention confusions, it is still hard to imagine what exactly it intends as being uniquely identified by '*a group of exactly three syringes*'. If it were three specific syringes in front of a specific physician here and now, then it would be a certain concrete aggregate of physical things. Since, however, the determiner is intended by HL7 to comprehend *any* combination of three syringes, it is not at all clear what, in reality, is intended—and certainly not at all clear how by these means we are advancing towards an adequate treatment of references to instances in healthcare communication.

8.4. Referent tracking versus information modelling

In fact, of course, 'entity,' in the context of HL7, does not mean 'entity' at all. For where referent tracking is indeed a matter of the tracking of entities in reality, HL7 subscribes rather to the paradigm of *information modelling*, which is a matter of the tracking (or modelling, or representation) of *information*. A referent tracking system can, as we have seen, track pieces of information about an entity (for example in the form of statements or images, which are then acknowledged as entities in their own right and are at the same time clearly distinguished from the entities about which they carry information). But unlike HL7 it is thereby able to distinguish in a clear and simple manner such pieces of information from the entities in reality which they are about.

The exclusive focus on information exhibited by HL7 yields, in contrast, an often painful awkwardness when attempts are made within its framework to refer also to the

corresponding entities in reality. Consider, for example, HL7's notion of '*mood*,' which is a determiner for acts used to convey continuity in the case where we need to communicate information about something like a physiological test. One mood asserts that an order for the test is made, another that the test is scheduled, another that the patient or specimen has arrived, another that the test is being conducted, another that the test report is generated, another that the report has been delivered, another that the report has been read, another that the test has been acted upon, and so forth. Although this looks like a sophisticated way of linking instances of different acts together in such a way as to preserve their common focus, it represents what is in fact a violation of sound ontological principles. When, for instance, a test is ordered or scheduled, then the test as such does not exist and there is no guarantee that it will ever do so: the patient might die, he might refuse the test, it may for some reason become inappropriate to perform the test, and so forth. Only in certain cases, therefore, is the test as ordered or scheduled properly to be represented as belonging to the *Entity* class (given that for HL7 '*an entity is a physical object that has, had or will have existence*') Certainly the '*plan*' to do such and such a test exists, but a plan to perform a test is not a kind, or 'mood,' of test. Rather it is another entity altogether—and this is so even when the test is at some later stage performed. Referent tracking, in contrast, since it tracks only entities in reality, must perforce take *existence* seriously, and it thus incorporates formal mechanisms to take account of this requirement (for example in the form of strict procedures for correcting erroneous entries).

The information-model paradigm is of course advanced also in other circles, and the fact that it leads to similar problems elsewhere suggests that there might be further reasons to pursue the realist paradigm, if only as a means of warding off the confusion of use and mention and counterintuitive treatments of the notion of existence. Consider, for example, the influential paper [49] of Rector et al., which contains assertions such as: '*Every occurrence level statement concerning the Jane Smith's Fracture of the Femur is an observation of the corresponding individual*'; whereby: '*The existence* [sic!] *of the individual Jane Smith's Fracture of Femur does not imply that Jane Smith has, or has ever had, a fracture of the femur* [sic!], *but merely that some observation has been made about Jane Smith regarding a fracture of the femur.*' [50].

This is not to deny that much valuable work has been invested in information model- and concept system-based tools for medical informatics. But we believe that the referent tracking paradigm must be called in aid to support any application of such tools at the point where EHRs and existing clinical terminologies come together. For the latter need to be applied to real cases in spatiotemporal reality, and to all that goes together therewith on the level of instances—to patients, their disorders, the particular treatments used. Referent tracking then gives us a means for allowing reality itself to serve as benchmark for such application [10], rather than the (more or less partial, more or less consistent) information about these entities that is packaged in this or that (more or less problematic) format.

In this respect, the referent tracking paradigm might contribute especially to the efforts of the TermInfo project set up by HL7 [51], which tries to use information models to define the interface to reality for the EHR. For the shallow and often self-contradictory semantics of existing HL7 models [52] currently leave it open whether the issues on TermInfo's agenda [53] could be resolved without some means for incorporating sound ontological principles into its general approach.

9. Conclusion

We have sketched, in very broad outline, how the referent tracking paradigm might be implemented in the healthcare environment, particularly in relation to clinical record-keeping. The key idea is to do full justice to the *what it is on the side of the patient* that is documented in an EHR. The effort that would be required to achieve this end from those involved in the documentation and date-entry process is, we believe, not significantly greater than is currently the case for those who are accustomed to working with concept-based systems. For rather than using such systems for retrieving *codes* that are applicable generically to the patient's condition, they would be used instead to retrieve the *IUIs* referring specifically to that concrete condition itself (just as we use proper names, and not general terms like 'human being,' to refer to individual patients). The invested effort will serve the direct purpose of providing better patient descriptions, and it will have the indirect consequence of leading to mappings between and to quality improvements within concept systems.

We foresee a time when, in addition to, or in replacement of, concept-based systems, interoperable principles-based ontologies will be widely used, incorporating formally defined relations that are appropriate for describing the linkages that obtain between particulars themselves [6]. From that point on there will become available to biomedical researchers a much richer description of real-world phenomena, of a type that will, we believe, be capable of being used for informatics-based biomedical research in a variety of still only barely imaginable ways. And again: the effort required to document relations between particulars by using ontologies of this sort will be not greater than that which is involved when using the unprincipled and at best poorly defined relationships that are found in the standard concept-based systems currently in use.

Acknowledgments

The present paper was written under the auspices of the Wolfgang Paul Program of the Alexander von Humboldt Foundation, the Volkswagen Project "Forms of Life," and the Network of Excellence in Semantic Interoperability and Data Mining in Biomedicine of the European Union.

References

[1] Popescu-Belis A, Lalanne D. Reference resolution over a restricted domain: references to documents. In Proceedings of the ACL2004 workshop on reference resolution and its applications, Barcelona, Spain; July 2004. p. 71–8.
[2] Ceusters W, Smith B, Kumar A, Dhaen C. Ontology-based error detection in SNOMED-CT. Proc MEDINFO 2004:482–6.
[3] Zanstra P, Rector A, Ceusters W, de Vries Robbé P. Coding systems and classifications in healthcare: the link to the record. Int J Med Inform 1998;48:103–9.
[4] Ceusters W. Language, medical terminologies and structured electronic patient records: how to escape the Bermuda Triangle. In De Moor G, De Clercq E, editors. Proc MIC; 2000. p. 7–14.
[5] Smith B, Ceusters W. Towards industrial-strength philosophy. How analytical ontology can help medical informatics. Interdisciplinary Sci Rev 2003;28(2):106–11.
[6] Smith B, Ceusters W, Klagges B, Kohler J, Kumar A, Lomax J, et al. Relations in biomedical ontologies. Gen Biol 2005;6(5):46.
[7] Ceusters W, Smith B. Referent tracking in electronic healthcare records. In: Engelbrecht R, et al., editors. Medical Informatics Europe. Amsterdam: IOS Press; 2005. p. 71–6.
[8] Smith B. Beyond concepts, or: ontology as reality representation. In: Varzi AC, Vieu L, editors. Formal ontology and information systems. Amsterdam: IOS Press; 2004. p. 73–84.
[9] Neuhaus F, Grenon P, Smith B. A formal theory of substances, qualities, and universals. In: Varzi AC, Vieu L, editors. Formal ontology in information systems. Amsterdam: IOS Press; 2004. p. 49–59.
[10] Smith B. Reasoning with biomedical terminologies. J Biomed Inform, in press.
[11] Willliams S, Kindel C. The component object model: A technical overview. October 1994. http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dncomg/html/msdn_comppr.asp [accessecd on April 12, 2005].
[12] ISO/IEC FDIS 9834-8:2004. Information technology—Open Systems Interconnection—Procedures for the operation of OSI Registration Authorities: Generation and registration of Universally Unique Identifiers (UUIDs) and their use as ASN.1 Object Identifier components.
[13] CCITT. Specification of basic encoding rules for Abstract Syntax Notation One (ASN.1), CCITT Recommendation X.209; January 1988.
[14] Bellare M, Rogaway P. The exact security of digital signatures—how to sign with RSA and Rabin. In: Proc Eurocrypt '96. LNCS 1070. Berlin: Springer; 1996. p. 399–416.
[15] Frege G. On sense and reference. In: Black M, Geach PT, editors. Translations from the philosophical writings of Gottlob Frege. Oxford: Blackwell; 1984.
[16] Netter WJ. Curing the unique health identifier: a reconciliation of new technology and privacy rights. Jurimetrics 2003;43(2):165–86.
[17] Wagner JC, Solomon WD, Michel PA, Juge C, Baud RH, Rector AL, et al. Multilingual natural language generation as part of a medical terminology server. In: Greenes RA et al., editors. Proc MEDINFO '95. Amsterdam: North-Holland; 1995. p. 100–4.
[18] Smith B. On the cognition of states of affairs. In: Mulligan K, editor. Speech act and sachverhalt: Reinach and the foundations of realist phenomenology. Dordrecht, Boston and Lancaster: Nijhoff; 1987. p. 189–225. Available: http://ontology.buffalo.edu/smith/articles/cogsvh/cogsvh.html .
[19] Clark T, Martin S, Liefeld T. Globally distributed object identification for biological knowledgebases. Brief Bioinform. 2004;5(1):59–70.
[20] Dogac A, Laleci G, Kabak Y, Unal S, Beale T, Heard S, Elkin P, Najmi F, Mattocks C, Webber D. Exploiting ebXML Registry Semantic Constructs for handling archetype metadata in healthcare informatics, Int J Metadata, Semant Ontol [forthcoming].
[21] ENV 13606-1 1999 Health informatics—electronic healthcare record communication— Part 1: extended architecture.
[22] Ceusters W, Waagmeester A, De Moor G. Syntactic–semantic tagging conventions for a medical treebank: the CASSANDRA approach. In: van derLei J, Beckers WPA, Ceusters W, van Overbeeke JJ, editors. Proc MIC'97. The Netherlands: Veldhoven; 1997. p. 183–1930.
[23] Ceusters W, Deville G. A mixed syntactic–semantic grammar for the analysis of neurosurgical procedure reports: the Multi-TALE experience. In Sevens C, De Moor G, editors. In Proc MIC '96; 1996. p. 59–68.
[24] Ceusters W, Rogers J, Consorti F, Rossi-Mori A. Syntactic-semantic tagging as a mediator between linguistic representations and formal models: an exercise in linking SNOMED to GALEN. Artif Intell Med 1999;15:5–23.
[25] Griffith SM, Kalra D, Lloyd DS, Ingram D. A portable communicative architecture for electronic healthcare records: the Good European Healthcare Record project Aim project A2014. Medinfo 1995;8(Pt1):223–6.
[26] Beale T, Heard S, Kalra D, Lloyd D. The openEHR EHR Information Model, version 4.5, December 10, 2004. http://www.openehr.org/repositories/spec-dev/latest/publishing/index.html. Accessed on April 18, 2005.
[27] Davis S. Tagging along. RFID helps hospitals track assets and people. Health Facil Manage 2004;17(12):20–4.
[28] Locke M. UC Considers Using Barcodes for Cadavers. San Francisco Chronicles, February 5, 2005. http://www.sfgate.com/cgi-bin/article.cgi?file=/news/archive/2005/02/04/national/a110118S59.DTL [accessed on April 19, 2005].
[29] Grenon P. Spatio-temporality in basic formal ontology: SNAP and SPAN, upper-level ontology, and framework for formalization. IFOMIS Technical Report, 05, 2003. http://www.uni-leipzig.de/~pgrenon/Downloads/grenon-tr1-part1.pdf [accessed on April 13, 2005].
[30] Rosse C, Mejino Jr JL. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. J Biomed Inform 2003;36(6):478–500.
[31] OBO relations ontology. http://obo.sourceforge.net/relationship/ [accessed on August 9, 2005].
[32] Donnelly M, Bittner T, Rosse C. Formal theory for spatial representation and reasoning in biomedical ontologies. Artif Intell Med [in press].
[33] Schneider L, Cunningham J. Ontological foundations of natural language communication in multi-agent systems. In: Palade V, Howlett RJ, Jain L, editors. Knowledge based intelligent information and engineering systems (LNAI-2773). Berlin: Springer; 2003. p. 1403–10.
[34] Gruber TR. A translation approach to portable ontology specifications. Know Acquisit 1993;5:199–220.
[35] Simon J, Fielding JM, Santos MDC, Smith B. Reference ontologies for biomedical ontology integration and natural language processing. Int J Med Inform, in press.
[36] Bodenreider O, Smith B, Burgun A. The ontology-epistemology divide: A case study in medical terminology. In: Varzi AC, Vieu L, editors. Proceedings of the third international conference on formal ontology in information systems (FOIS 2004). IOS Press; 2004. p. 185–95.
[37] Ceusters W, Smith B. A terminological and ontological analysis of the NCI Thesaurus. Methods Inform Med [in press].
[38] Ceusters W, Smith B, Kumar A, Dhaen C. Ontology-based error detection in SNOMED-CT. Proc MEDINFO 2004:482–6.
[39] ISO-1087-1:2000. Vocabulary of terminology.
[40] ISO-1087:1990. Vocabulary of terminology.
[41] Martucci D, Masseroli M, Pinciroli F. Gene ontology application to genomic functional annotation, statistical analysis and knowledge mining. Stud Health Technol Inform 2004;102:108–31.
[42] Medicomp Systems I: MEDCIN. Chantilly, VA: Medicomp Systems; 1998.

[43] Weed L. Medical records that guide and teach. N Engl J Med 1968;278:593–600.
[44] Barrows RC, Johnson SB. A data model that captures clinical reasoning about patient problems. In: Gardner RM, editor. Proceedings of the 19th Annual Symposium Computer Applications in Medical Care; 1995. p. 402–505.
[45] Huff SM, Rocha RA, Bray BE, Warner HR, Haug PJ. An event model of medical information representation. J Am Med Inform Assoc 1995;2:116–34.
[46] Coyle JF, Mori AR, Huff SM. Standards for detailed clinical models as the basis for medical data exchange and decision support. Int J Med Inform 2003;69(2-3):157–74.
[47] HL7 Version 3 Standard, January 2005 Ballot. http://www.hl7.org/ [accessed on August 1, 2005].
[48] Smith B., Ceusters W, Temmerman R. Wüsteria. In Proc MIE 2005. p. 647–652.
[49] Rector AL, Nowlan WA, Kay S, Goble CA, Howkins TJ. A framework for modelling the electronic medical record. Methods Inform Med 1993;32(2):109–19.
[50] Rector AL, Nowlan WA, Kay S, Goble CA, Howkins TJ. A framework for modelling the electronic medical record. Methods Inform Med 1993;32(2):109–19.
[51] http://www.terminfo.org [accessed on August 2, 2005].
[52] Vizenor L. Actions in health care organizations: an ontological analysis. Medinfo 2004;11(Pt 2):1403–7.
[53] NASA summit on composition and context. Houston, August 2–3, 2004. http://www.tc215wg3.nhs.uk/pages/pdf/wg3_179.pdf [accessed on August 2, 2005].