Introducing Realist Ontology for the Representation of Adverse Events

Werner CEUSTERS [a,1], Maria CAPOLUPO [b], Georges DE MOOR [c], Jos DEVLIES [c]
[a] *New York State Center of Excellence in Bioinformatics & Life Sciences, University at Buffalo, NY, USA*
[b] *Independent consultant, Buffalo NY, USA*
[c] *Research in Advanced Medical Informatics and Telematics VZW, Ghent, Belgium.*

Abstract. The goal of the REMINE project is to build a high performance prediction, detection and monitoring platform for managing Risks against Patient Safety (RAPS). Part of the work involves developing an ontology enabling computer-assisted RAPS decision support on the basis of the disease history of a patient as documented in a hospital information system. A requirement of the ontology is to contain a representation for what is commonly referred to by the term '*adverse event*', one challenge being that distinct authoritative sources define this term in different and context-dependent ways. The presence of some common ground in all definitions is, however, obvious. Using the analytical principles underlying Basic Formal Ontology and Referent Tracking, both developed in the tradition of philosophical realism, we propose a formal representation of this common ground which combines a reference ontology consisting exclusively of representations of universals and an application ontology which consists representations of defined classes. We argue that what in most cases is referred to by means of the term '*adverse event*' – when used generically – is a *defined class* rather than a *universal*. In favour of the conception of adverse events as forming a defined class are the arguments that (1) there is no definition for '*adverse event*' that carves out a collection of particulars which constitutes the extension of a universal, and (2) the majority of definitions require adverse events to be (variably) the result of some observation, assessment or (absence of) expectation, thereby giving these entities a nominal or epistemological flavour.

Keywords. Basic Formal Ontology, Referent Tracking, adverse events, patient safety.

1. Introduction

'*High performance prediction, detection and monitoring platform for patient safety risk management (REMINE)*' is the name of a European Large Scale Integrating Project (IP) which has been funded by the European Commission since January 1, 2008 [1]. The main objective is to develop a technological platform and identify best practice business processes allowing automated management and prevention of Risks against Patient Safety (RAPS).

Part of the work to be carried out consists of the development of an ontology that will support several functionalities offered by the envisioned technological platform. In this

[1] Corresponding author: Werner Ceusters, Ontology Research Group, NYS Center of Excellence in Bioinformatics & Life Sciences, 701 Ellicott street, suite B2-160, Buffalo NY 14203, USA.

Table 1. Adverse event related definitions from authoritative sources

ID	Term	Definition	Source	Ref.
D1	***adverse drug event*** (adverse drug error)	Any incident in which the use of a medication (drug or biologic) at any dose, a medical device, or a special nutritional product (for example, dietary supplement, infant formula, medical food) may have resulted in an adverse outcome in a patient.	JTC	[2]
D2	***adverse drug experience***	any adverse event associated with the use of a drug in humans, whether or not considered drug related, including the following: • an adverse event occurring in the course of the use of a drug product in professional practice; • an adverse event occurring from drug overdose whether accidental or intentional; • an adverse event occurring from drug abuse; • an adverse event occurring from drug withdrawal; and • any failure of expected pharmacological action.	FDA	[3]
D3	***adverse drug reaction***	an undesirable response associated with use of a drug that either compromises therapeutic efficacy, enhances toxicity, or both.	JTC	[2]
D4	***adverse event***	an ***observation*** of a change in the state of a subject ***assessed*** as being untoward by one or more ***interested parties*** within the ***context*** of a protocol-driven research or public health.	BRIDG	[4]
D5	***adverse event***	an event that results in unintended harm to the patient by an act of commission or omission rather than by the underlying disease or condition of the patient	IOM	[5]
D6	***adverse event***	any unfavourable and unintended sign (including an abnormal laboratory finding), symptom, or disease temporally associated with the use of a medical treatment or procedure that may or may not be considered related to the medical treatment or procedure	NCI	[6]
D7	***adverse event***	any untoward medical occurrence in a patient or clinical investigation subject administered a pharmaceutical product and which does not necessarily have to have a causal relationship with this treatment	CDISC	[7]
D8	***adverse event***	an untoward, undesirable, and usually unanticipated event, such as death of a patient, an employee, or a visitor in a health care organization. Incidents such as patient falls or improper administration of medications are also considered adverse events even if there is no permanent effect on the patient.	JTC	[2]
D9	***adverse event***	an injury that was caused by medical management and that results in measurable disability.	QUIC	[8]
D10	***error of omission***	An error which occurs as a result of an action not taken. Errors of omission may or may not lead to adverse outcomes.	JTC	[2]
D11	***observation***	an act of recognizing and noting a fact or an occurrence of an event of interest. An observation may involve examination, interviews, or measurement with devices. Observations are not intended to alter the state of the subject.	BRIDG	[4]
D12	***serious adverse drug experience***	Any adverse drug experience occurring at any dose that results in any of the following outcomes: • death • a life-threatening adverse drug experience • inpatient hospitalization • prolongation of existing hospitalization • a persistent or significant disability/incapacity • a congenital anomaly/birth defect	FDA	[3] [9]

paper, we focus on one particular RAPS issue, namely the ontological treatment of *adverse events*.

The term '*adverse event*' is defined in the literature in a variety of ways, superordinate terms frequently used being '*reaction*', '*effect*', '*event*', '*problem*', '*experience*', '*injury*', '*symptom*', '*illness*', '*occurrence*', '*change*', and even '*something*', '*act*', '*observation*' and '*term*', the latter four being the result of applying flawed terminological theories which rest on a confusion between an entity and an observation or record thereof [10]. This multitude of definitions is brought about by the many organisations and initiatives that have set themselves the noble goal of reducing the occurrence of adverse events, especially since the year 2000, when the Institute of Medicine published its report *To Err is Human: Building a Safer Health System* [11]. **Table 1** contains a small selection of adverse event definitions by authoritative sources, drawn from a larger collection that we composed for our work in [12].

Research aimed at bringing about some order in this domain falls into three categories. One is classification, as witnessed by the work of Chang *et al*. who developed – on the basis of a set of criteria specified in [13] – a classification schema consisting of five root nodes which they found to be the '*homogeneous elements*' encountered in relevant sources: ***Impact***, ***Type, Domain, Cause*** and ***Prevention and Mitigation*** [14]. Others, such as the BRIDG consortium, have tried to resolve the multitude of definitions by reaching consensus on just one [4], with the result of being extremely reductionist. A third group of researchers has focused on building ontologies. Unfortunately, this latter group has typically employed the rather weak principles underlying the '*concept*'-orientation [15] in ontology development, so that, for example, '*age*' and '*gender*' become a subclass of '*patient*' [16]. We nonetheless believe that ontology is indeed the right approach to take in addressing this difficult and important problem, but, in contrast to what is still the majority view among ontologists, we believe that ontology will bring benefits only when rigorous principles are applied, principles that go far beyond the basic requirement of computational soundness.

2. Objective and design

Our goal is to bring clarity to the terminological wilderness that grew out of all the efforts documented in [12]. Problems arise not only because of differences amongst initiatives in terms of scope, health care settings involved, jurisdictions, and objectives – the consequence being that definitions resulting from such efforts are not applicable outside the original boundaries – but also because of a widespread failure to adopt sound ontological and terminological principles in analysing and conveying what is relevant. As an example, a definition such as " *'Adverse outcome' should be understood to mean not only a non-trivial adverse outcome* […] *but also an incident* […] *which results in a recognized potential risk of a non-trivial adverse outcome* […]" [17] (irrelevant detail omitted), is of the form '*an X is an X or a Y which leads to an X*' and is thus at best uninformative.

To obtain our goal, we analysed the literature and collected all relevant definitions and descriptions that we found. We modified some of these definitions slightly in order to have them convey better what we judged to be the intended message, thereby still keeping track of the original versions in order to identify general principles for improved definition construction in the domain of patient safety.

We also studied a variety of classification systems, taxonomies, terminologies and concept-based ontologies – we use the term '*concept-based ontologies*' to differentiate representational artifacts created on the basis of the concept orientation from the realism-based ontology that is being developed under REMINE – in order to obtain a comprehensive list of entity types whose nature and interrelationships are to be studied and formally represented to satisfy requirements of the project.

3. Methodology

We performed our analysis following the principles advocated in Basic Formal Ontology and Referent Tracking.

3.1. Basic Formal Ontology

Basic Formal Ontology (BFO) is a framework encapsulating best practices in ontology development that is designed to serve as basis for the creation of high-quality shared ontologies in the biomedical domain [18, 19].

BFO has a *realist* orientation based on the view that terminologies and ontologies are to be aligned not on '*concepts'* but rather on entities in reality [15]. Central to this view are three assumptions.

The first is that biological reality exists objectively in and of itself, i.e. independently of the perceptions or beliefs or theories of cognitive beings. Thus not only do a wide variety of entities exist in reality (human beings, hearts, bacteria, disorders, ...), but also how these entities relate to each other – that certain hearts are parts of human beings, that certain bacteria cause disorders in human beings – is not a matter of agreements made by scientists but rather of objective fact.

The second assumption is that reality is accessible to us and that its structure can be discovered: it is scientific research that allows human beings to find out what entities exist and what relationships obtain between them.

The third assumption is that an ontology should mirror its corresponding domain of reality. Thus an important aspect of the quality of an ontology is determined by the degree to which not merely its individual representational units correspond to entities in reality but also the structure according to which these units are organized mimics the corresponding structure of reality. Realism-based ontology development was introduced into biomedical informatics some fifteen years ago as a means of detecting and avoiding the systematic mistakes characteristic of concept-based terminologies [20-24], mistakes which are not eliminated through the use of description logics or similar computational devices [25].

BFO acknowledges only those entities which exist in reality, and rejects all types of putative entities postulated merely as artifacts of specific logical or computational frameworks. The corresponding logical and computational artifacts themselves, however, are indeed accepted as part of reality. BFO captures a small number of basic categories into which the entities in reality are divided, thereby distinguishing, at the highest level of its organisation:

(1) *particulars* such as ***Werner Ceusters*** from *universals* such as HUMAN BEING,[2]

(2) *continuants* such as ***Werner Ceusters' heart*** from *occurrents* such as ***the beating of Werner Ceusters' heart***,

(3) *independent entities* such as ***Werner Ceusters' heart*** from *dependent entities* such as ***the function of Werner Ceusters' heart***.

Dependent entities are such that they cannot exist – in the ontological rather than biological sense – without some instance of the category independent entity. BFO also distinguishes three major families of relations between entities in the categories just distinguished:

(1) *<p, p>-relations*: from particular to particular (for example: ***Werner Ceusters' s brain*** being part of ***Werner Ceusters***);

(2) *<p, u>-relations*: from particular to universal (for example: ***Werner Ceusters*** being an instance of HUMAN BEING);

(3) *<u, u>-relations*: from universal to universal (for example: HUMAN BEING being a subkind of ORGANISM) [26].

3.2. Referent Tracking

Referent tracking has been introduced as a new paradigm for entry and retrieval of data in the Electronic Health Record (EHR) to avoid the multiple ambiguities that arise when statements in an EHR refer to disorders, lesions and other entities on the side of the patient exclusively by means of generic terms from a terminology or ontology [27]. Referent tracking avoids such ambiguities by introducing *IUIs* – Instance Unique Identifiers – for each numerically distinct entity that exists in reality and that is referred to in statements in a record. As ontologies serve integration of information at the level of the types (universals and defined classes) which particulars instantiate, so referent tracking serves integration of information at the level of these particulars themselves, which, if they are catered for at all in current record systems, are represented in heterogeneous and unstable ways.

Drawing on this framework, we have proposed a calculus for use in quality assurance of the complex representations created for clinical or research purposes, for example in coding of clinical trial data [28]. This calculus is based on a distinction between three levels [29]:

(1) the level of reality (for example, in the medical domain, the reality on the side of the patient);

(2) the cognitive representations of this reality for example as embodied in observations and interpretations on the part of clinicians and others;

(3) the publicly accessible concretizations of these cognitive representations in artifacts of various sorts, of which ontologies and terminologies and Electronic Health Records are examples.

---

[2] For clarity, we will from here on represent particulars in ***bold italic*** and universals in SMALL CAPS. Terms (or other representational units) denoting either universals or particulars will be written in italics between single quotes. For additional clarity, we will sometimes use the words '*particular*', '*universal*' and '*term*' explicitly to denote entities of the corresponding type.

4. Results

4.1. Terminological conventions

In line with the terminology proposed in [29], we will henceforth use the term '*class*' to denote a collection of all and only those particulars to which a given general term applies. A class can be either: (1) the extension of a universal, thus comprehending all and only those particulars which instantiate the corresponding universal (at that time); or (2) a subset of the extension of a universal defined as being such that the *members* of this class exhibit an additional property which is (a) not shared by all instances of the universal, and (b) also (can be) exhibited by particulars which are not instances of that universal. For such a class, we reserve the term '*defined class*'. Examples are: the class of influenza patients in Leipzig; the class of rabbits with congenitally absent nipples.

We will further use the term '*property*' to denote the combination formed by a relation that this particular enjoys with some other entity and this entity itself. For example, it is a property of my brain that it is part_of me; the property here is the combination formed by the part_of relation and my body. Similarly, it is a property of my brain that it is an instance of BRAIN; the property here is the combination formed by the instance_of relation and the universal BRAIN. Note that such combinations are not extra entities which exist in addition to the entities or relations through which they are formed. Rather, to talk of properties is to parse the reality already existing in the context of a given particular in a new way, reflecting the subject-predicate structure of languages such as English and the 'F(a)' structure of predicate logic based languages.

By '*portion of reality*' we mean any combination of particulars (including classes and defined classes), universals and properties. The use of this expression, too, does not reflect any extra entities which would exist in addition to the entities or relations which already exist and are classified under other headings.

We use the term '*representational unit*' (RU) for any symbolic representation (a code, a character string, an icon, …) which *denotes* a portion of reality and which is not constructed out of smaller parts which play a similar denoting role.**Table 2** gives an overview of the type of representational artifacts that are useful for representing portions of reality and of the sorts of entities that should be represented in each type of artifact. The latter is inspired by the view that *reference ontologies* should be the equivalent of scientific theories and therefore should represent what is generic in the world – whether or not in a specific domain – in a way that maximizes faithfulness and comprehensiveness with respect to reality.

Table 2: representational artifacts and their suggested representational units

Representational artifact	Contains representational units for …
Reference Ontology	• universals • relationships between universals (following the principles of the Relation Ontology [26])
Application Ontology	• universals • defined classes • relationships between universals and defined classes (following the principles of the Relation Ontology [26]) • particulars required for defining defined classes
Inventory	• particulars • properties

Application ontologies, in contrast represent matters in a local, purpose-driven way and in a format that is more suitable for computation [30]. Examples of *inventories* are databases which store information about particulars, examples being Electronic Health Records or Adverse Event Registries.

4.2. Core representational units

Table 3 shows the minimal collection of classes related to entities in reality that must be taken into consideration if we are to be in a position to represent the portion of reality around a particular patient on whose side an adverse event might have occurred under any of the definitions for adverse event analyzed thus far. Under the label '*denotation'* we propose a generic term applicable to a member of the corresponding class. The '*Class type*' column indicates whether the class is the extension of a universal (U) or a defined class (DC). The '*Particular type*' column indicates to what category of particulars, in terms of Basic Formal Ontology, the members of the corresponding class belong.

The descriptions provided in the right-most column are, be it noted, not to be interpreted as definitions for the terms that we choose to use in our ontology to denote the corresponding entities. Rather, they serve only to illustrate the sorts of roles played by different sorts of entities in a scenario in which an adverse event might have occurred. It is important, too, that the terms listed under the denotation-column should be seen as pertaining to the domain of adverse events. Thus for example we do not claim that anything which would be referred to by third parties by means of the term '*observation*' falls under the description provided. The conditionals that are used in most of these descriptions reflect the fact that a particular portion of reality might be such that a phenomenon which is considered to be an adverse event under one definition, is not an adverse event in terms of another definition. The conditionals should not be interpreted as having in every case to do with probabilities or uncertainty.

4.3. The place of 'adverse events'

The representational units for the core classes identified above can be used to represent all possible portions of reality which feature entities that can be referred to by means of the term '*adverse event*' under any of the definitions listed in [12]. As an example, **Table 4** lists the particulars and associated properties involved in a case in which

a patient born at time t_0

undergoing anti-inflammatory treatment and physiotherapy since t_2

for an arthrosis present since t_1

develops a stomach ulcer at t_3.

This table thereby provides an example of an adverse event case analysis of the sort that is made possible by the framework here presented.

The relationships employed in composing representations of properties in this Table are drawn from [26, 31]. We preserve the formatting conventions proposed in [26], except that we pick out particulars using ***bold italic***. We introduce the primitive **is_about** relation holding between a representational unit and the entity in reality about which this unit contains information at a certain time. We further take certain shortcuts in our representation of the temporal relationships involved in such an analysis, by simply stating for example that t_0 **earlier** t_1 **earlier** t_2 **earlier** t_3.

Table 3: Universals and Defined Classes for the adverse events domain.

	Denotation	Class Type	Particular Type	Description (role in adverse event scenario)
				Level 1
C 1	subject of care	DC	independent continuant	person to whom ***harm*** might have been done through an ***act under scrutiny***
C 2	act under scrutiny	DC	act of care	***act of care*** that might have caused ***harm*** to the ***subject of care***
C 3	act of care	U	process	activity carried out by a ***care giver*** to a ***subject of care***, motivated by an ***underlying disease*** and a ***care intention***
C 4	care giver	DC	independent continuant	person that performed an ***act of care*** directed to the ***subject of care***
C 5	underlying disease	DC	dependent continuant	the disease in the ***subject of care*** which is part of what serves to motivate performance of the ***act of care***
C 6	involved structure	DC	independent continuant	anatomical structure (of the ***subject of care***) involved in an ***act of care***
C 7	structure change	U	process	change in an anatomical structure of a person
C 8	structure integrity	U	dependent continuant	aspect of an anatomical structure deviation from which would bring it about that the anatomical structure would either (1) itself become dysfunctional or (2) cause dysfunction in another anatomical structure
C 9	integrity change	U	structure change	change in the ***structure integrity*** bringing about a change in the range of circumstances under which the anatomical structure would become dysfunctional or cause dysfunction in another structure
C 10	harm	U	integrity change	***integrity change*** bringing about an expansion in the range of circumstances of the sort typically occurring in the life of the ***subject of care*** under which the anatomical structure would become dysfunctional or cause dysfunction in another structure
C 11	care effect	DC	integrity change	***integrity change*** brought about by an ***act of care***
C 12	subject investigation	DC	process	looking for a ***structure change*** in the ***subject of care***
C 13	harm assessment	U	process	determining whether an ***observation*** is faithful to reality, and if so, whether the ***structure change*** which is the target of the ***observation*** is a ***harm***
C 14	care intention	DC	dependent continuant	intention of a ***care giver*** that motivates him towards an ***act of care***
				Level 2
C 15	observation	DC	dependent continuant	cognitive representation of a ***structure change*** resulting from an act of perception within a ***subject investigation***
C 16	harm diagnosis	DC	dependent continuant	cognitive representation, resulting from a ***harm assessment***, and involving an assertion to the effect that a ***structure change*** is or is not a ***harm***
C 17	care effect belief	DC	dependent continuant	belief on the side of the ***care giver*** concerning the ***care effect*** that he ascribes to the ***act of care***
				Level 3
C 18	care reference	DC	information entity	concretized (through text, diagram, …) piece of knowledge drawn from state of the art principles that can be used to support the appropriateness of (or correctness with which) processes are performed involving a ***subject of care***

Table 4: Example of an adverse event case analysis

IUI	Particular description	Properties
#1	the patient who is treated	***#1*** **member** C1 **since** t_2
#2	***#1***'s treatment	***#2*** **instance_of** C3 ***#2*** **has_participant** ***#1*** **since** t_2 ***#2*** **has_agent** ***#3*** **since** t_2
#3	the physician responsible for #2	***#3*** **member** C4 **since** t_2
#4	***#1***'s arthrosis	***#4*** **member** C5 **since** t_1
#5	***#1***'s anti-inflammatory treatment	***#5*** **part_of** ***#2*** ***#5*** **member** C2 **since** t_3
#6	***#1***'s physiotherapy	***#6*** **part_of** ***#2***
#7	***#1***'s stomach	***#7*** **member** C6 **since** t_2
#8	***#7***'s structure integrity	***#8*** **instance_of** C8 **since** t_0 ***#8*** **inheres_in** ***#7*** **since** t_0
#9	***#1***'s stomach ulcer	***#9*** **part_of** ***#7*** **since** t_3
#10	coming into existence of ***#9***	***#10*** **has_participant** ***#9*** **at** t_3
#11	change brought about by ***#9***	***#11*** **has_agent** ***#9*** **since** t_3 ***#11*** **has_participant** ***#8*** **since** t_3 ***#11*** **instance_of** C10 **at** t_3
#12	noticing the presence of ***#9***	***#12*** **has_participant** ***#9*** **at** t_{3+x} ***#12*** **has_agent** ***#3*** **at** t_{3+x}
#13	cognitive representation in ***#3*** about ***#9***	***#13*** **is_about** ***#9*** **since** t_{3+x}

We also allow for temporal annotations additional to those described in [26], at the same time remaining faithful to EN 12388:2005: Health Informatics – Time Standards for Healthcare Specific Problems [32].

Under the proposed scenario, ***#10***, i.e. the appearance of ***#9***, would (modulo the wide variation in interpretations that can be given to the majority of the definitions found) qualify as an adverse event as defined by the Institute of Medicine (definition D5).

However, for definition D9, it would rather be ***#9*** itself that would so qualify, while for D4, the definition of 'adverse event' proposed by the BRIDG consortium [4], it would be either ***#12*** or ***#13***. The counterintuitive nature of the latter case has its roots in certain conflations in the HL7 RIM [33], by which BRIDG is heavily inspired.

Because of the various sorts of entities that qualify as adverse events depending on which definition is used, at least two adverse event classes need to be defined: one for adverse events under views that see adverse events as processes, and one for adverse events that see them as continuants. A further distinction has to be made between adverse events as entities in first order reality, and phenomena in first order reality qualified as adverse events by relating to certain cognitive representations, records or theories.

5. Discussion

Already a very superficial analysis of the definitions in **Table 1** applying the analytical principles just sketched demonstrates that the question "*What are adverse events?*" cannot be answered directly, but needs to be reformulated as "*What might the author of a particular sentence containing the phrase 'adverse event' be referring to when he uses that phrase?*". Indeed, the authors of the listed definitions must have had very distinct entities in mind: we cannot imagine even one single example of an entity which would be such that, were it placed before these authors, they would each in turn be able

to point to it while the first would say – faithfully and honestly – "*that is an observation*" (definition D4), the second: "*that is an injury*" (definition D9), the third: "*that is a laboratory finding*" (definition D6), and so on. Clearly, nothing which **is** an injury can **be** a laboratory finding, although, of course, laboratory findings can aid in diagnosing an injury or in monitoring its evolution. Similarly, nothing which **is** a laboratory finding, can **be** an observation, although, of course, some observation must have been made (by either a human being or a device) if we are to arrive at a laboratory finding.

One could argue, perhaps, that the authors of some of these definitions resort to metonymy, i.e. linguistic formulations in which a term denoting some entity is replaced by a term that denotes a related entity as in 'The White House decided that …', rather than 'The President of the United States decided that …'. If that would be the case, we would still have to qualify such usage as bad practice, specifically because we are convinced that definitions should be constructed to avoid ambiguity, rather than to contribute to confusion. This is all the more the case where the definitions in question are to serve as the basis for reasoning systems developed for use by computers.

However, because all the authors of the mentioned systems use the term '*adverse event*' in some context for a variety of distinct entities, and because these contexts look quite similar – in each of them, more or less the same sort of entities seem to be involved – there is some common ground (some portion of reality) which is such that the entities within it can be used as referents for the various meanings of '*adverse event*'.

5.1. Classifying adverse event related entities in terms of the three levels of reality

The definitions for the term '*adverse event*' and for other closely related terms differ amongst themselves in that they require a representation which resorts to one, two or all three levels of reality as described above. The first part of D12 (from the Food and Drug Administration) is an example in which all terms refer to level 1 entities: *drugs*, *drug doses*, *deaths*, *hospitalizations*, *disabilities*, and so forth, are all entities that exist in first-order reality. Another example is D9: the terms '*injury*', '*medical management*', '*measurement*' and '*disability*', when used in the context of a specific patient that may or may not have experienced an adverse event, all denote existing entities on the side of that particular patient and his environment, and are not about something else: these terms thus denote level 1 entities. D2, in contrast, requires bringing level 2 and perhaps even level 3 entities into the picture, and this because of the clause '*any failure of <u>expected</u> pharmacological action*'. Expectations can only be raised by a cognitive being and are part of the cognitive representation this cognitive being has constructed ***about*** the first order reality which forms his environment. Thus, in this interpretation of D2, i.e. if the expectation concerning the pharmacological action is 'in the mind' of the particular clinician assessing whether the patient has an adverse drug experience, D2 involves a level 2 entity. However, if this expectation is something which is part of '*general knowledge*' or belongs to the '*state of the art*', then we are dealing with an additional level 3 entity: in order for the clinician assessing the case to have access to that '*general knowledge*', it must have been concretized in some enduring fashion, for example in a manual or textbook.

5.2. Lack of clarity in definitions

D2 exhibits a characteristic which, unfortunately, is shared by the majority of the definitions encountered: they lack sufficient clarity of phrasing to allow an analysis to be conducted unproblematically in realist terms. Often multiple interpretations can be given to one or more terms used within such a definition, whereby each interpretation suggests a denotation at a distinct level of reality. An example is definition D3, in which the response that is described as being *undesirable* can be understood in three different ways:

(1) as denoting something on level 1, namely a *realizable entity* (a *disposition* or *tendency* [34]), which exists objectively as an increased health risk; in this sense any event '*that either compromises therapeutic efficacy, enhances toxicity, or both*' is undesirable;
(2) as denoting something on level 2, so that, amongst all of those events which influence therapeutic efficacy or toxicity, only some are considered *undesirable* (for whatever reason) by either the patient, the caregiver or both; or
(3) as denoting something relating to level 3, so a particular event occurring on level 1 is *undesirable* only when it is an instance of a type of event that is listed in some guideline, good practice management handbook, i.e. in some published statement of the state of the art in relevant matters.

In other cases, this sort of analysis results in detecting hidden assumptions, conflations or even serious inconsistencies either within one definition or in the combination of several definitions offered by the same source.

An example of an inconsistency within one single definition when the latter is analyzed in realist terms is provided by the attempt at a literal interpretation of D5, and more precisely of the use, there, of the term '*act of omission*', especially if, as suggested by D10, that term is taken in such a way that it does not denote anything which exists either now or in the past. In Referent Tracking terms, there would thus be nothing to which a IUI could be assigned. Indeed, while we believe that the phrase 'action not taken' is a linguistic description (level 3 entity) that can be used adequately and meaningfully in reporting some feature of a complex portion of reality (level 1 entity), such a use does not yet signify that the term denotes directly some entity in that portion of reality. While terms of the form 'doing something' do have referents in first order reality, there are no such referents denoted by terms like 'doing nothing'.

Consider the example given in [5], in which '*not testing a diabetic patient for* HbA_{1c}' is stated to be an '*act of omission*'. This is because, in result of the work of the Diabetes Quality Improvement Project [35, 36], it is considered bad practice not to do such a test at regular intervals [37]. But clearly, if all that exist as relevant first order entities are a patient's disease (here, the diabetes) and some adverse event , then it is not possible that some '*act of omission*' – i.e. some *not doing something* that one is supposed to do according to the state of the art – could be the cause of the adverse event. The only such cause would here be the underlying disease. Events, so we believe, can only be caused by what exists. And it is in the given case indeed clear that it is precisely the diabetes on the side of the patient that causes the adverse event, although it is true that, if the test had been taken, along with further appropriate actions in line with the results of that test, then it could be expected that no adverse event would have occurred. Therefore, a better definition for what D5 is trying to express would be: '*an event that results in unintended harm to the patient (1) through an act of commission*

rather than through some underlying disease or condition of the patient, or (2) through an underlying disease or condition of the patient in the absence of appropriate actions which should have been taken in line with the state of the art in dealing with the disease'. This rephrased definition accounts better for something else that the Institute of Medicine almost certainly had in mind when producing D5, namely that many acts of commission are part of a procedure which, in order to be conducted *lege artis*, must include taking actions of a sort which, if they would not be taken, would lead to harm to the patient because of the act of commission. An example is incising an artery during some surgical procedure in a way that inevitably leads to bleeding. It would be inappropriate, in such case, not to take actions to reduce the bleeding. Here it is not the underlying disease which leads to harm to the patient, and nor is it the 'not stopping the bleeding' which leads to the harm. Rather it is the bleeding caused by the incision.

6. Conclusion

We have used the principles of Basic Formal Ontology (BFO), including the Relation Ontology (RO), and Referent Tracking (RT) as an analytical framework to study the ontological nature of what is denoted by the term '*adverse event*'. Our research indicates that this framework is adequate to serve a number of important purposes, and that, when used appropriately, it avoids the inconsistencies and incompatibilities inherent in other approaches. Nevertheless, some further developments, especially in RO are required if we are to be able to deal more formally with some extensions that we proposed here: (1) a family of relations to deal with various aspects of *aboutness* and *denotation* to relate level 2 and level 3 entities to level 1 entities, (2) a *membership* relation to link particulars to defined classes, and (3) the capacity to refer to (open-ended) time periods in addition to time instants.

7. Acknowledgments

We wish to thank Barry Smith for useful comments on our ontological analysis.

8. References

[1] REMINE Consortium. REMINE: Supporting Hospitals in Risk Management. 2008 [cited 2008 April 9]; Available from: http://www.remine-project.eu/

[2] The Joint Commission. Sentinel Event Glossary of Terms. 2008 2008 [cited 2008 April 14]; Available from: http://www.jointcommission.org/SentinelEvents/se_glossary.htm

[3] U.S. Food and Drug Administration. IND Safety Reports. In: Code of Federal Regulations Title 21 Volume 5 § 312.32, ed. 2007.

[4] Weng C, Becich M, Fridsma D. Collective Domain Modeling across Clinical Trials Standards: Needs, Challenges, and Design Implications. *Information Technology and Communications in Health (ITCH)*. Victoria, BC, Canada 2007.

[5] Committee for Data Standards for Patient Safety - Institute of Medicine. Patient Safety: Achieving a New Standard of Care. Washington DC: The National Academies Press 2004.

[6] National Cancer Institute. Common Terminology Criteria for Adverse Events (CTCAE) Version 3.0. 2006.

[7] CDISC Submission Data Standards Team. CDISC SDTM Implementation Guide (Version 3.1.2) - Study Data Tabulation Model Implementation Guide: Human Clinical Trials; 2007 July 25.

[8] Quality Interagency Coordination Task Force. Glossary of Terms. 2008 [cited 2008 April 14]; Available from: http://www.quic.gov/report/mederr8.htm
[9] U.S. Food and Drug Administration. Code of Federal Regulations; Title 21--Food And Drugs; Chapter I--Food And Drug Administration; Department Of Health And Human Services ; Subchapter D--Drugs For Human Use; Part 310 -- New Drugs; Subpart D--Records and Reports. 2007.
[10] Smith B, Ceusters W, Temmerman R. Wüsteria. In: Engelbrecht R, Geissbuhler A, Lovis C, Mihalas G, eds. *Connecting Medical Informatics and Bio-Informatics Medical Informatics Europe 2005*. Amsterdam: IOS Press 2005:647-52.
[11] Institute of Medicine. To Err is Human: Building a Safer Health System. Washington DC: National Academy Press 2000.
[12] Ceusters W, Capolupo M, Moor GD, Devlies J. Various Views on Adverse Events: a collection of definitions. 2008 April 20, 2008 [cited 2008 April 29]; Available from: http://www.org.buffalo.edu/RTU/papers/AdverseEventDefs.pdf
[13] The WHO World Health Organization Alliance for Patient Safety. Project to Develop the International Patient Safety Event Taxonomy: Updated Review of the Literature 2003-2005; Final Report. Geneva: World Health Organization; 2005 September.
[14] Chang A, Schyve PM, Croteau RJ, O'leary DS, Loeb JM. The JCAHO patient safety event taxonomy: a standardized terminology and classification schema for near misses and adverse events. International Journal for Quality in Health Care Advance Access. 2005;17(2):95-105.
[15] Smith B. From Concepts to Clinical Reality: An Essay on the Benchmarking of Biomedical Terminologies. Journal of Biomedical Informatics. 2006;39(3):288-98.
[16] Jeong S, Kim H-G. Design of Semantically Interoperable Adverse Event Reporting Framework. *The Semantic Web – ASWC 2006*. Berlin / Heidelberg: Springer 2006:588-94.
[17] College of physicians and surgeons of Newfoundland and Labrador. Disclosure of an Adverse Outcome. 2008 [cited 2008 April 16]; Available from: http://www.cpsnl.ca/PolicyDocument.asp?ID=21
[18] Spear AD. Ontology for the Twenty First Century: An Introduction with Recommendations. Saarbrucken, Germany: Institute for Formal Ontology and Medical Information Science 2006.
[19] Smith B, Ashburner M, Ceusters W, Goldberg L, Mungall C, Shah N, et al. The OBO Foundry: Remolding Biomedical Ontologies to Support Data Integration. Nature Biotechnology. 2007;25:1251-5.
[20] Johansson I. Bioinformatics and Biological Reality. Journal of Biomedical Informatics. 2006;39(3):274-87.
[21] Ceusters W, Smith B, Goldberg L. A terminological and ontological analysis of the NCI Thesaurus. Methods of Information in Medicine. 2005;44:498-507.
[22] Ceusters W, Smith B, Kumar A, Dhaen C. Ontology-based error detection in SNOMED-CT®. In: Fieschi M, Coiera E, Li Y-CJ, eds. *MEDINFO 2004*. Amsterdam, The Netherlands: IOS Press 2004:482-6.
[23] Ceusters W, Smith B, Kumar A, Dhaen C. Mistakes in medical ontologies: Where do they come from and how can they be detected? In: Pisanelli DM, ed. *Ontologies in Medicine Studies in Health Technology and Informatics*. Amsterdam, The Netherlands: IOS Press 2004:145-64.
[24] Smith B. Ontology and the Logistic Analysis of Reality. In: Guarino N, Poli R, editors. Proceedings of the International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation; 1993; Padova: Institute for Systems Theory and Biomedical Engineering of the Italian National Research Council; 1993. p. 51-68.
[25] Ceusters W, Smith B. Ontology and Medical Terminology: why Descriptions Logics are not enough. *Towards an Electronic Patient Record (TEPR 2003)*. San Antonio 2003.
[26] Smith B, Ceusters W, Klagges B, Köhler J, Kumar A, Lomax J, et al. Relations in biomedical ontologies. Genome Biology. 2005;6(5):R46.
[27] Ceusters W, Smith B. Strategies for Referent Tracking in Electronic Health Records. Journal of Biomedical Informatics. 2006 June;39(3):362-78.
[28] Ceusters W, Smith B. A Realism-Based Approach to the Evolution of Biomedical Ontologies. *Proceedings of AMIA 2006* 2006:121-5.
[29] Smith B, Kusnierczyk W, Schober D, Ceusters W. Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain. *KR-MED 2006, Biomedical Ontology in Action*. Baltimore MD, USA 2006.
[30] Menzel C. Reference Ontologies — Application Ontologies: Either/Or or Both/And? 2007 [cited 2008 April 28]; Available from: http://bioontology.org/wiki/images/d/d9/MenzelOntology.pdf
[31] Smith B, Grenon P. The Cornucopia of Formal-Ontological Relations. Dialectica. 2004;58(3):279-96.
[32] European Committee for Standardization. EN 12388:2005. Health informatics - Time standards for healthcare specific problems. 2005.

[33] Smith B, Ceusters W. HL7 RIM: An Incoherent Standard. In: Hasman A, Haux R, Lei Jvd, Clercq ED, Roger-France F, eds. *Studies in Health Technology and Informatics Ubiquity: Technologies for Better Health in Aging Societies - Proceedings of MIE2006*. Amsterdam: IOS Press 2006:133-8.
[34] Jansen L. Tendencies and other Realizables in Medical Information Sciences. The Monist. 2007 October 2007:(in press).
[35] McLaughlin S. The Diabetes Quality Improvement Project. Diabetes Spectrum. 2000;13(1):5-10.
[36] American Diabetes Association. Standards of Medical Care in Diabetes. Diabetes Care. 2005;28:S4-S36.
[37] National Committee for Quality Assurance. HEDIS 2008 Technical Specifications for Physician Measurement National Committee for Quality Assurance 2007.