ReMINE

High performances prediction, detection and monitoring platform for patient safety risk management

FP7 Contract: 216134

– Deliverable –

D4.2 – RAPS Domain Ontology (M12 Version)

Document Information

Document Name:	ReMINE_D4.2_RAMIT_WP4_080601
Revision:	V011
Revision date:	January 11, 2009
Author:	Werner Ceusters, Maria Capolupo, Jos Devlies (for RAMIT VZW)
Contributors:	
Security:	Public

Diffusion list

Consortium, European Commission, general public

Approvals

	Name	Beneficiary	Date	Visa
Project Coordinator	Michele CARENINI	GMD		
Technical Manager				

Document history

Revision	Date	Modification	Author
V01	Dec 15, 2008	Prefinal draft	RAMIT
V011	Jan 11 2009	Final document	RAMIT

Table of Content

1 Executive Summary

In this deliverable we describe the background materials and methodology used to develop the RAPS Domain Ontology, the latter itself being still work in progress in line with the developments in the course of the ReMINE project. At this stage, the work performed includes

(1) a representational framework for describing faithfully what is the case in reality in the context of adverse events, consisting of adequate representational units for the relevant entities in reality, and a language for expressing how the entities in reality that are denoted by these units relate to each other;

(2) a classification method and system that allows descriptions of adverse events to be categorized under distinct views; and

(3) a terminology in which terms denote unambiguously (a) entities or relationships in reality, (b) representational units used in the framework, and (c) classes from the classification system.

The work presented here does not – for reasons explained in section 4.1 – follow the *concept-based* approach in ontology development, but, in contrast, a *realist* agenda which expands considerably the possibilities of the former. The realist agenda is based on the assumptions that (1) reality exists objectively in itself, i.e. independent of the perceptions or beliefs of cognitive beings; (2) reality, including its structure, is accessible to us, and can be discovered through (scientific) research; and (3) the quality of an ontology is at least determined by the accuracy with which its structure mimics the pre-existing structure of reality.

After a short introduction covering the ReMINE project and the role of ontology therein (chapter 2), we present a requirements analysis in chapter 3 which includes a number of background materials covering work on adverse events that our final version of the ontology should be compatible with. This is followed by an explanation of the general methodology and approach in chapter 4. Our methodology builds further on the well-established principles adhered to in Basic Formal Ontology (chapter 5) and Referent Tracking and the application thereof to the Disease Ontology (chapter 6) which provides a general framework for clinically relevant entities such as diseases, disorders, diagnoses, signs, symptoms, and so forth. We restrict the scope however to that what is relevant in the context of adverse events, thereby keeping pace with the needs of the pilot sites of ReMINE.

In chapter 7 we present a detailed analysis of the confusing ways in which adverse events are dealt with in the literature and offer a representational framework that can accommodate all views found thus far. This provides the introduction to chapter 8 which contains a semi-formal description of the upper part of the RAPS Domain Ontology which will continue to evolve and grow in line with future developments in ReMINE, as well as with the relevant ontologies which are being developed as part of the OBO-Foundry initiative (addressed in chapter 4) as explained in chapter 9.

2 Introduction

'High performances prediction, detection and monitoring platform for patient safety risk management (ReMINE)' is a European Large Scale Integrating Project (IP) funded by the European Commission under contract 216134 [1]. The statement of work was approved 03 December 2007, and work started Jan 1, 2008.

The two main objectives of the ReMINE project are:

(1) to develop a new technological platform that is able to perform semi-automated RAPS (Risks Against Patient Safety) management, and

(2) to propose organisational changes with considerable added value in relevant environments.

The first objective is being achieved by developing the ReMINE system, a RAPS identification and analysis system for the acquisition and mining of relevant multimedia data present in hospitals. This system will then be used to predict, detect and monitor RAPS related events in the collaborating facilities.

Tasks to achieve the second objective include developing adequate clinical risk management processes, establishing a more active role for RAPS managers, and identifying new ways for interactions amongst different health care professionals in a local health care system to solve RASP issues.

An essential component of the ReMINE system is an ontology with associated taxonomy and terminology that will support several functionalities offered by the envisioned technological platform. Table 1 summarises the initial desiderata for this ontology.

Three ReMINE deliverables cover these components. D4.1 describes the domain of adverse events from a cognitive perspective, as perceived by clinicians, adverse event and risk managers, and so forth, i.e. human beings. This description uses the machinery of DOLCE which has been specifically developed – as witnessed by its own title: Descriptive Ontology for Linguistic and Cognitive Engineering – to be able to deal with cognitive issues. The result of the effort described in D4.1 is a taxonomy which provides the "model of use" with which clinicians are familiar.

D4.2 and D4.3 do not look exclusively at the adverse event domain from a cognitive human perspective, but from a perspective that can be understood by machines, including how the domain is cognitively perceived by humans. Indeed, ontologies developed for machine-understanding need to take into account several additional levels of detail that humans can deal with implicitly, but machines can not. The use of Basic Formal Ontology for the work described in these deliverables, in contrast to DOLCE, makes it possible to have a smooth integration with the cognitive perspective by the recognition of three levels of reality in BFO of which the cognitive realm is only one.

Further in line with the state of the art, we make a distinction between a *domain ontology* and an *application ontology*. Whereas the former is intended to be a purpose-independent representation of the portion of reality covered by a domain, the latter is a derivation of the former in light of a specific application.

This deliverable covers the RAPS Domain Ontology, whereas deliverable D4.3 describes the RAPS Application Ontology.

Table 1. Initial ReMINE desiderata for an ontology for managing risks against patient safety.

A. Domain coverage	A1: adverse events; A2: medical errors; A3: information needs and information seeking behaviour; A4: communication errors.
B. Ontology language requirements	B1: interpretable by software agents responsible for adverse event detection, risk classification, guideline execution and the identification of correlation between data and a DSS engine to allow users to perform simulations and "What if" scenario analysis; B2: usable within the Federated Enterprise Reference Architecture (FERA) framework [2]; B3: OWL compatible.
C. Interoperability requirements	C1: hospital information systems; C2: relevant guidelines; C3: CDC's Public Health Conceptual Data Model [3]; C4: patient safety taxonomies developed by WHO [4-6], DATIX [7] and NRLS [8]

3 Requirements analysis

3.1 '*Adverse event*' in the biomedical literature

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 [9]. 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* [10]. **Table 2** contains a small selection of adverse event definitions by authoritative sources, drawn from a larger collection that we composed for our work in [11].

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 [5] – 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*** [12]. Others, such as the BRIDG consortium, have tried to resolve the multitude of definitions by reaching consensus on just one [13], 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 [14] in ontology development, so that, for example, '*age*' and '*gender*' become a subclass of '*patient*' [15]. 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.

3.2 ReMINE's cognitive engineering notion of '*adverse event*'

Also within the ReMINE project, it was originally not specified what is denoted by the term 'adverse event': the Description of Work does not contain a definition. In Deliverable D4.1 [16] a cognitive engineering position is defended to the effect that an adverse event is:

- an 'incident [that] occurred during the past and [is] documented in a database of adverse events' (p23);
- a 'perdurant' (p26);
- 'that occurs to a patient' (p23);
- an expectation of some future happening that can be prevented (p23).

Although each assertion above contains useful information that can easily be understood by human beings, it is not possible to formalise the information in exactly the same way in order to have it unambiguously understood by machines, and this for several reasons. First, there is no such thing in reality that exhibits all the stated properties at once, or that can acquire or loose them throughout its existence. Clearly, nothing which happened in the past (an occurrent), can be an expectation (a continuant).

Second, it is not explicitly specified what sorts of incidents are considered to be adverse events.

Table 2**. 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	[17]
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	[18]
D3	***adverse drug reaction***	an undesirable response associated with use of a drug that either compromises therapeutic efficacy, enhances toxicity, or both.	JTC	[17]
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	[13]
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	[19]
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	[20]
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	[21]
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	[17]
D9	***adverse event***	an injury that was caused by medical management and that results in measurable disability.	QUIC	[22]
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	[17]
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	[13]
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	[18] [23]

Under one reading, a machine (or software agent running in a machine) could assume that an incident becomes an adverse event by the mere fact of the event being reported 'in a database of adverse events'. But that – if applied cautiously – would violate the principle that the past cannot be changed: something which is not an adverse event at the time it happens, can not become one afterwards!

This cognitive engineering view leaves also many questions unanswered:

- Who has the authority to add reports?
- What criteria are used by that authority?
- How to deal with false positives and negatives?

If ReMINE partners want to add reports on incidents to the adverse event database (and they do so as specified in [16], p22), then further criteria must be established for what counts as adverse event. A decision to this effect must for sure be taken by the risk managers who will use the ReMINE system in the future.

3.3 Materials related to the patient safety domain

3.3.1. JCAHO Patient Safety Event Taxonomy

The lack of a common language in the US national discussions on patient safety that followed the publication of the 2003 Institute of Medicine report '*Patient Safety: Achieving a New Standard of Care*' [24], motivated the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) to develop a common terminology and classification schema for collecting and organizing patient safety data [12]. The project comprised (1) a systematic literature review; (2) evaluation of existing patient safety terminologies and classifications to identify those that should be included in the core set of a standardized taxonomy; (3) assessment of the taxonomy's face and content validity; (4) the gathering of input from patient safety stakeholders in multiple disciplines; and (5) a preliminary study of the taxonomy's comparative reliability. Using principles that were earlier proposed by the World Health Organization (WHO) [4, 5], JCAHO identified five '*root nodes*' for the classification:

- ***Impact***: the outcome or effects of medical error and systems failure, commonly referred to as harm to the patient;
- ***Type***: the implied or visible processes that were faulty or failed;
- ***Domain***: the characteristics of the setting in which an incident occurred and the type of individuals involved;
- ***Cause***: the factors and agents that led to an incident;
- ***Prevention and Mitigation***: the measures taken or proposed to reduce incidence and effects of adverse occurrences.

The root nodes were then further divided into 21 subclassifications which in turn were subdivided into more than 200 coded categories and an indefinite number of uncoded text fields to capture narrative information.

The taxonomy was later used as input for WHO's *International Classification for Patient Safety* [6].

3.3.2. SHEL model

The SHEL model, introduced in 1972 in aviation, describes the interrelationships and interdependencies of different system components on the one hand and human factors on the other

hand under the influence of environmental conditions [25]. It has since then been applied to document and analyze the occurrence of errors in a variety of domains, including healthcare [26].

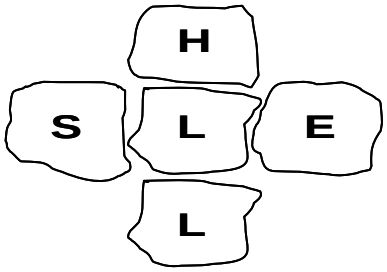

Figure 1: SHEL - model

The model obtained its name from the four main components of systems as perceived by researchers in Human Factors, a discipline consisting of several areas of research including human performance, technology design, and human-computer interaction:

- ***S***oftware: procedures, protocols, training,…;
- ***H***ardware: machines, medical instruments, …;
- ***E***nvironment: operating theatres, wards, consultation rooms, laboratories;
- ***L***iveware: human factors, health care professionals, patients, ….

The uneven edges of the system components suggest that they are constantly changing and will never match perfectly. Errors, therefore, occur not only inside components, but also at the various interfaces, examples being erroneous communications (L-L), non-adherence to protocols (L-S), manipulation mistakes (L-H), stress-inducing environments (L-E).

3.3.3. The CDC's Public Health Conceptual Data Model

The Centers for Disease Control and Prevention (CDC) published the first version of the Public Health Conceptual Data Model in 2000 [3]. The purpose of the model is to document the information needs of public health so that the CDC and its state and local partners in public health can (1) establish data standards for public health, (2) collaborate with national health informatics standards setting bodies to

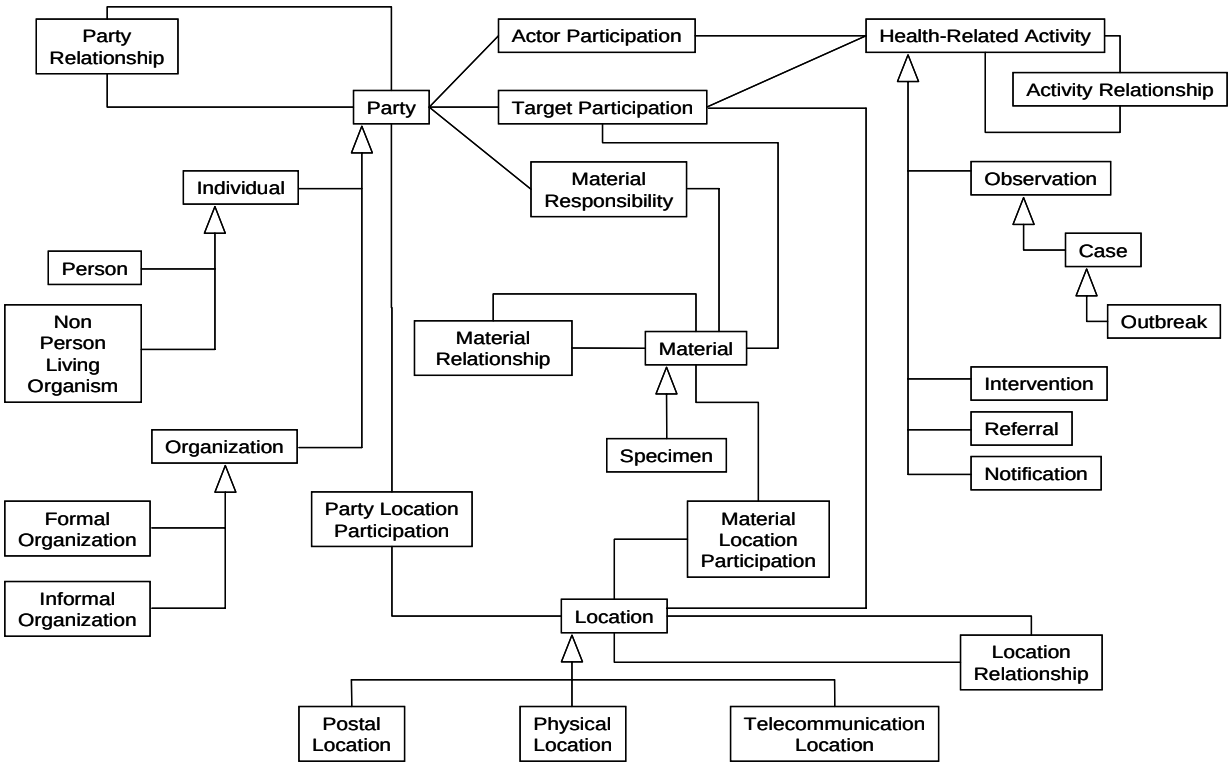

Figure 2: CDC's Public Health Conceptual Data Model Diagram

define standards for the exchange of information among public health agencies, and healthcare providers, and (3) construct computerized information systems that conform to established standards for use in the management of data relevant to public health. It further supports the CDC's National Electronic Disease Surveillance System (NEDSS), a collection of complementary computerized information systems designed to automate the process of gathering health data, facilitate the monitoring of the health of communities, assist in the analysis of trends and detection of emerging public health problems, and provide information for setting public health policy.

3.3.4. The International Classification for Patient Safety

The WHO launched in October 2004 the World Alliance for Patient Safety in response to a World Health Assembly Resolution urging WHO and Member States to pay the closest possible attention to the problem of patient safety [27]. The Alliance decided to create the International Classification for Patient Safety (ICPS) [5] in such a way that it would be applicable across the full spectrum of healthcare from primary care to highly specialized areas and usable in conjunction with existing processes and systems [6].

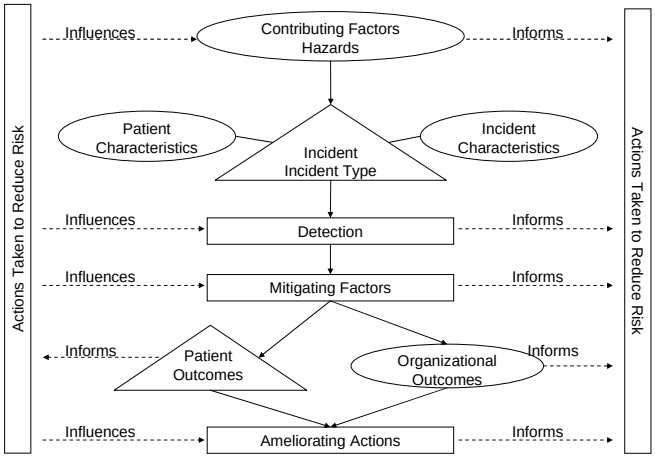

Figure 3: Conceptual Framework for the International Classification for Patient Safety

The conceptual model shown in Figure 3 serves as a reference for the definition of 46 core terms (Table 3). The majority of these terms serve in turn as indexes to schemas that allow a detailed classification of the reasons, circumstances, persons, and so forth, involved in a patient safety issue, for example, '*incident type – clinical administration – process – response to emergency*'.

Table 3: Core terms of the International Classification for Patient Safety

Accountable	Error	Patient outcome
Actions taken	Event	Patient safety
Adverse event	Harm	Patient safety incidents
Adverse reaction	Hazard	Preventable
Agent	Health	Quality
Ameliorating action	Healthcare	Resilience
Attributes	Healthcare-associated harm	Risk
Circumstance	Incident characteristics	Root cause analysis
Class	Incident type	Safety
Classification	Injury	Semantic relationship
Concept	Mitigating factor	Side effect
Contributing factor	Near miss	Suffering
Degree of harm	Organizational outcome	System failure
Detection	Patient	System improvement
Disability	Patient characteristics	Violation
Disease		

3.3.5. AHRQ's Common Formats

The US *Patient Safety and Quality Improvement Act of 2005* establishes a framework by which healthcare providers may voluntarily report information regarding patient safety events and quality of care through Patient Safety Organizations (PSOs) [28]. August 2008, the Agency for Healthcare Research and Quality (AHRQ) published *The Common Formats* which contain definitions for data elements and reporting formats that allow healthcare providers to collect and submit standardized information regarding patient safety events to PSOs [29, 30]. Information about three types of events are hoped to be collected through this voluntary initiative:

(1) incidents, defined as patient safety events that reached the patient, whether or not there was harm,

(2) near misses (or close calls), i.e. patient safety events that did not reach the patient, and

(3) unsafe conditions, defined as any circumstance that increases the probability of a patient safety event.

The Common Formats were designed by an interagency Patient Safety Work Group that aligned the Common Formats, 'to the extent practicable' [30], with the ICPS. The list of entities recognised can be found in chapter 10.

3.4 RAPS taxonomy

Within the ReMINE project, early versions of the RAPS *taxonomy* are developed prior to the RAPS *ontology*: several pilot sites, each active in a specific medical discipline, contribute by submitting terms that are relevant for patient safety risk management in their specialty. Terms are either (1) *general,* covering aspects of risk management that are applicable in every context, an example being '*technical failure*', or (2) *specific*, thus in the context of situations that are only encountered in specific medical disciplines, such as, for instance, the term '*primigravida over 39*' which is used in obstetrics and gynaecology. The terms are then organized into a taxonomy by the ReMINE partner responsible for this task. This methodology is motivated by the fact that it is easier and less time consuming to collect terms from the domain of discourse and to classify them according to some conceptual schema, than to perform a detailed analysis of the entities in reality that is covered by these terms. Furthermore, whereas the purpose of the RAPS ontology is to enable automatic decision support, the taxonomy must function as interface between user and application. This requires the organization of the terms not as much to reflect how the domain is objectively structured, but rather how users think and communicate about it. The taxonomy is based upon the ReMINE model as shown in Figure 4.

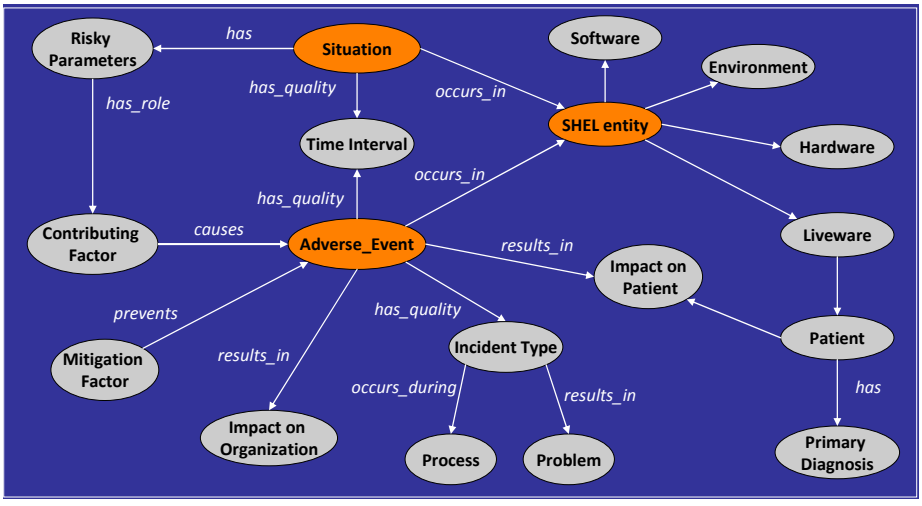

Figure 4: ReMINE model (adapted from [16])

The June 2008 version of the RAPS taxonomy contained 366 terms whereas the August 2008 version 644. Specific terms were drawn from the domain of Obstetrics and Gynaecology whereas the general terms were borrowed from the JCAHO classification [12] and the SHEL-model [26] which both served also as inspiration for the taxonomic backbone, all of them classified under the three major categories of the DOLCE ontology [31]. Table 4 shows the top three levels of the taxonomy.

Table 4: Top three levels of the June 2008 version of the ReMINE Taxonomy

Endurant	Contributing Factor	Environment contributing factor
		Hardware contributing factor
		Liveware contributing factor
		Software contributing factor
	Impact on Organization	
	Impact on Patient	
	Mitigation Factor	Environment mitigation factor
		Hardware mitigation factor
		Liveware mitigation factor
		Software mitigation factor
	Primary Diagnosis	
	Problem	
	Risky parameters	
	SHEL entity	Environment
		Hardware
		Liveware
		Software
Perdurant	Adverse Event	
	Process	
	Situation	
Quality	Degree of Harm	
	Incident Type	
	Patient Quality	
	Time Interval	

The RAPS Taxonomy uses the upper level divisions of the DOLCE ontology for its own structure, such that the RAPS Taxonomy terms are classified under the top-level entities *endurant*, *perdurant* and *quality* as perceived by cognitive beings. A similar top-level structure is used in Basic Formal Ontology (BFO) and in many cases, a direct mapping is possible: DOLCE's endurants correspond with BFO's independent continuants, perdurants are roughly similar to BFO's processual entities, whereas some of DOLCE's qualities are similar to BFO's qualities which are all dependent continuants rather than enjoying a top level category on their own. However, there are cases in which the cognitive engineering view differs from the realism-based view: some entities, for example, that are perceived as perdurants in the RAPS Taxonomy, are continuants under the BFO perspective and vice versa. In these cases, a more complex mapping, often taking the form of bridging axioms, is required.

3.5 Preliminary analysis

The definitions found in the literature and the objectives of the ReMINE project suggest that there are at least five different entities – which we further refer to using the identifiers #1, #2, and so forth – that can qualify for being an adverse event:

- #1: an incident that happened in the past;
- #2: the interpretation by some cognitive agent that #1 is an adverse event;
- #3: the expectation by some cognitive agent that similar incidents might happen in the future;
- #4: an entry in the adverse event database concerning #1;

- #5: an entry in some other system about #3 for mitigation or prevention purposes.

Thus if an incident (#1) ***is judged to be*** an adverse event, there is some #2 in, for instance, a clinician about #1 which for that reason – applying ReMINE's definition for adverse event – requires the introduction of #4 in a database. It is further because of #3 that also #5 should be generated because if indeed #1 ***is*** an adverse event, similar incidents may happen in the future and mitigation efforts are required.

Of course, judgments may be wrong, Some possibilities are:

- presence of #1 with unjustified absence of #2:
  - #1 was not perceived at all, or
  - #1 was not assessed as being an adverse event
- unjustified presence of #2:
  - there was no #1 at all, or
  - #1 was not an adverse event
- unjustified absence of #4:
  - for the same reasons as under the first bullet above, or
  - justified presence of #2 but no report about it in the database.

A complete analysis of the possible combinations of (un)justified presence or absence of each of the entities listed above is shown in Table 5.

An additional complexity follows from the fact that distinct clinicians, depending on what definition they apply, may hold different beliefs about whether a specific incident such as #1 (1) really happened, (2) is of a specific sort, or (3) counts as an adverse event. They may further differ in beliefs about what caused the incident, and about how to prevent future happenings of incidents of the same sort. Moreover, they may change their beliefs over time.

Table 5: Registration of adverse events and mitigation factor: combinations of co-existence

	Past incident related			Mitigation related	
	AE happened	AE perception	AE database entry	Interpreted	Registered
Case	#1	#2	#4	#3	#5
1	+	-	-	-	-
2	+	+	-	-	-
3	+	+	-	+	-
4	+	+	-	+	+
5	+	+	+	-	-
6	+	+	+	+	-
7	+	+	+	+	+
8	-	-	-	-	-
9	-	+	-	-	-
10	-	+	-	+	-
11	-	+	-	+	+
12	-	+	+	-	-
13	-	+	+	+	-
14	-	+	+	+	+

Whether an incident is an adverse event (under one or more definitions adhering to the realist agenda) is a matter of objective fact, and not a matter of consensus. What are matters of consensus are (1) definitions for what should be counted as adverse events despite the fact that they can be applied wrongly and themselves possibly be in error, (2) policies about registration, and (3) policies about mitigation and prevention, although, whether they are effective, is again a matter of objective fact.

3.6 Conclusions

The preliminary analysis just presented provides arguments to justify a realism-based approach towards ontology and more specifically the use of the principles underlying Basic Formal Ontology (BFO) [32] and Referent Tracking [33] as explained in chapter 5.

The five types of entities (#1, #2, …) presented above indeed fit nicely in the 3-layered structure of reality as argued for in BFO: #1 and #2 are in first-order reality, #3 is a cognitive representation (level 2), and #4 and #5 are level-3 entities, i.e. representations – in contrast to cognitive representations which, as long as we can't read somebody's mind, are not accessible to other cognitive agents than the one entertaining the representation – that are accessible to cognitive agents that are observing and interacting in first-order reality

The need to learn from judgment mistakes made in the past, to identify differences in assessment skills, and to advance the state of the art by providing better evidence for identifying causes and consequences, developing better treatments with less iatrogenic effects or adopting better mitigation and prevention strategies, requires the RAPS domain and application ontologies to cover all three levels of reality, in addition to appropriate error management [34].

4 Methods and design

4.1 Fundamentals

Our approach is based on the ontology development strategy put forward by the OBO-Foundry, an endeavour that has been quite successful in biomedicine [35, 36].

The principles of the Foundry can be summarized, in their current version, as follows [37]:

- First, are syntactic principles to the effect that an ontology must employ one or another common shared syntax, possess a unique identifier space, and have procedures for identifying distinct successive versions.
- Second, are principles involving definitions:
  - textual definitions (and, by degrees, equivalent formal definitions) are to be provided for all terms;
  - terms and definitions must be composed using *the methodology of cross-products* which is the view that where ontologies need to include complex representations these should be built up compositionally out of component representations already defined within other, more basic feeder ontologies if available.
  - ontologies must use relations that are unambiguously defined according to the pattern set forth in the OBO Relation Ontology (RO) [38].
- Third, ontologies are required to be open, have a clearly specified and clearly delineated content, have a plurality of independent users, and be subject to a collaborative development process involving the developers of other ontologies covering neighbouring domains.
- Finally, the *principle of orthogonality* asserts that for each domain there should be convergence upon a single one.

The fourth principle, to be practical, requires to make a clean distinction between *reference ontologies* [39] and *application ontologies*.

Reference ontologies are analogous, although in different ways, to both scientific theories and textbooks. Each has its own subject-matter, which consists of the entities in reality addressed by the corresponding domain of science (hence the quasi-synonym *domain ontologies*). Each seeks to maximize descriptive adequacy to this subject-matter by being built out of representations which are correct when viewed in light of our best current scientific understanding and should exhibit the following features [37]:

(1) be a common resource that cannot be bought or sold,

(2) represents a well-demarcated scientific domain;

(3) is subject to constant maintenance by domain experts

(4) is designed to be used in tandem with other, complementary ontologies, and

(5) is independent of format and implementation.

Application ontologies, in contrast, are comparable to engineering artifacts. They are constructed for specific practical purposes such as RAPS management.

Sadly, however, the predominating view, primarily in circles of computer scientists and knowledge engineers, is that ***all*** ontologies are engineering and computer science artifacts which are nothing

more than 'just another application' of the developers' computational expertise, and thus as something that is of lesser scientific importance than core computer science issues for example in logic or in systems for ontology mapping [40]. The result has been that many ontologies and the terminologies that can be seen as their predecessors are full of mistakes [41-44] which are not eliminated – although often so argued – through the use of description logics or similar computational devices [45]. As further pointed out in [40], this '*self-limiting approach*' of the computer science approach of ontology design will in the end '*not be able to exploit the full potential of the ontology idea*', and the authors accordingly insist that the ontologies developed for scientific purposes need to be taken much more seriously as first-class citizens by computer scientists and knowledge engineers.

4.2 Approach

Based on these considerations, and the goals of the ReMINE project, our objective was to identify the core components of an ontology that

(1) would allow software agents to detect risks against patient safety,

(2) satisfies the A and C criteria of the ReMINE desiderata (Table 1) which includes the challenge of remaining forward compatible with the ReMINE taxonomy for risks against patient safety (*RAPS taxonomy*) as it evolves based on end-user requests from collaborating pilot sites, and

(3) follows the principles for high quality ontology design set forward by the Open Biomedical Ontology Foundry (OBO-Foundry) initiative [36].

To obtain our goal, we studied the model underlying the July and August 2008 versions of the RAPS taxonomy and compared it with other models for risk management that we obtained from the literature. We also studied a number of domain-related classification systems, taxonomies, terminologies and concept-based ontologies – we use the term '*concept-based ontologies*' to differentiate such representational artifacts clearly from the realism-based ontology that is developed under ReMINE – in order to obtain a comprehensive list of entity types whose nature and interrelationships were to be studied and formally represented to satisfy the ReMINE requirements.

We performed our analysis following the principles advocated in Basic Formal Ontology and Referent Tracking, the principles of which are explained in chapter 5.

5 Realism-based ontology

5.1 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 [36, 46]. According to BFO, an ontology is '***a representational artifact comprising a taxonomy as proper part, whose representational units are intended to designate some combination of universals, defined classes, and certain relations between them***' [47]. BFO follows a realist agenda in which a distinction between three levels of reality is recognized:

(1) the level of first-order 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.

BFO is a representational artifact (level 3) which is composed in modular fashion of sub-representations, the smallest such unit being a *representational unit*, describing primarily first-order entities (level 1).

5.1.1. BFO's categories

BFO captures a small number of basic categories into which reality is divided thereby distinguishing at the highest level of its organisation

(1) particulars from universals, the former being non-repeatable entities that can exist only in one place and during one period of time such as ***Werner Ceusters***, the latter being *that* in virtue of which a thing is what it is, and without which that thing would not be the kind of thing that it is, an example being HUMAN BEING,[1] and

(2) continuants from occurrents, the former being defined as entities that exist in full at any time in which they exist at all, persist through time while maintaining their identity, and having no temporal parts, an example being ***Werner Ceusters' heart*** and the latter defined as entities that have temporal parts and that happen, unfold or develop in time, an example being ***the beating of Werner Ceusters' heart***.

BFO recognizes three disjoint types of continuants:

(1) spatial regions, i.e. independent continuants that do not bear qualities nor inhere in any other entities, such as ***the space occupied by Werner Ceusters yesterday at noon***,

(2) independent continuants – defined as continuants that are the bearers of qualities and realizables, in which other entities inhere and which themselves cannot inhere in anything – examples being ***Werner Ceusters*** and ***Werner Ceusters' heart***, and

(3) dependent continuants – defined as continuants that inhere in or are born by other entities – of which ***Werner Ceusters's shape*** is an example.

An important type of dependent continuant in the context of risks against patient safety (and beyond) is that of a realizable entity. Realizable entities inhere in continuant entities and are not

[1] 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.

exhibited in full at every time in which they inhere. The realization of a realizable entity is a particular manifestation, functioning, or process that occurs under certain circumstances.

Table 6: Relevant formal relations for an ontology of risks against patient safety

Relatum	Relation	Relatum
continuant	genidentical-with	continuant
independent continuant	transformation-of	independent continuant
independent continuant	derivation-of	independent continuant
continuant	part-of	continuant
occurrent	part-of	occurrent
continuant	segmentation-of	occurrent
independent continuant	participates-in	processual entity
	perpetrates	
	initiates	
	perpetuate	
	terminates	
	influences	
	facilitates	
	hinders	
	prevents	
	mediates	
	patient-of	
dependent continuant	realizes	processual entity
dependent continuant	inheres-in	independent continuant
processual entity	involves	independent continuant
	creates	
	sustains-in-being	
	degrades	
	destructs	
	affects	
	demarcates	
	blurs	
processual entity	affects	dependent continuant
	creates	
	sustains-in-being	
	degrades	
	destructs	
continuant	located-in	continuant
continuant	contained-in	continuant

BFO also recognizes three disjoint types of occurrents:

(1) processual entities – defined as occurrents that exist in time by occurring or happening, have temporal parts, and always depend on at least one continuant – such as ***Werner Ceusters' life***,

(2) spatiotemporal regions – defined as occurrent entities at or in which processual entities can be located – such as ***the spatiotemporal region occupied by Werner Ceusters' life***, and

(3) temporal regions – defined as occurrent entities that are part of time – such as ***the time period during which Werner Ceusters wrote this paragraph***.

Importantly, BFO acknowledges only those entities which exist in biological reality, and rejects all those types of putative negative entities – absences, non-existents, possibilia, and the like – which are postulated merely as artifacts of specific logical or computational frameworks. This doesn't mean that BFO is not expressive enough to represent, for instance, that some patient is missing a left leg, but rather that the representation is built up exclusively of representational units that denote something that exists [48].

5.1.2. BFO's relationships

BFO distinguishes three major families of formal 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*** at time t); (2) <p, u>-relations: from particular to universal (for example: ***Werner Ceusters*** being an instance of HUMAN BEING at time t); and (3) <u, u>-relations: from universal to universal (for example: HUMAN BEING being a subkind of ORGANISM) [38].

Of the relationships that have been documented thus far [38, 49], Table 6 contains the relationships and the corresponding sort of entities that can figure as their relata deemed relevant for our purposes.

5.2 Referent Tracking

Referent Tracking (RT) is a paradigm for information management that is distinct from other approaches in that each data element has to point to a *portion of reality* (see further) in a number of predefined ways. Whereas BFO covers through its representation of universals what is generic, RT deals with what is specific, thus offering a method to describe particulars in terms of BFO, other realism-based ontologies, and even concept-based terminologies.

RT has been introduced 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 [33]. 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 the record. It's applicability has been studied in a variety of domains outside healthcare, including security and intelligence informatics [50], digital rights management [51], and corporate memories in knowledge management systems [52].

A referent tracking system has been developed to store and retrieve descriptions in the form of RT-tuples [53]. As shown in Figure 5, RT-tuples – see [50] for their abstract syntax and semantics – are descriptions that correspond to *configurations*, a configuration being a portion of reality in which a number of entities, the latter being anything that exists or has existed in the past, stand in certain relationships to each other. RT-tuples are representational constructs which are built out of representational units – linked together by means of relational expressions – which can be:

(1) IUIs, which uniquely and directly identify particulars,

(2) UUIs, which uniquely and indirectly identify universals through their representation in some realism-based ontology (the latter itself being identified in the RT-tuple by means of a IUI), or

(3) CUIs, which either indirectly identify collections of particulars or nothing at all. The latter is the case when a CUI corresponds to a concept (drawn from a concept-based terminology) that doesn't denote anything existing [14, 43, 48].

CUIs that do denote, denote one of the following:

(1) the *extension* of a universal, which consists of all the particulars that instantiate the universal (at a given time), for instance all hospitals;

(2) a *defined class*, which consists of a collection of particulars out of the extension of a universal that exhibit an additional property which is (a) not shared by all instances of the universal, and (b) also might be exhibited by particulars which are not instances of that universal, for instance all hospitals in New York State,

(3) an *ad hoc* (or *composite*) *class*, which is a collection of particulars not satisfying (1) or (2), for instance all hospitals and clinicians in New York State.

RT-tuples are not only used to describe configurations involving non-referring particulars using relationships as listed in Table 6, but also to describe how data, information, and knowledge relate to that what these information bearers are supposed to be about. This includes RT-tuples themselves, for instance to express that certain descriptions are wrong [54].

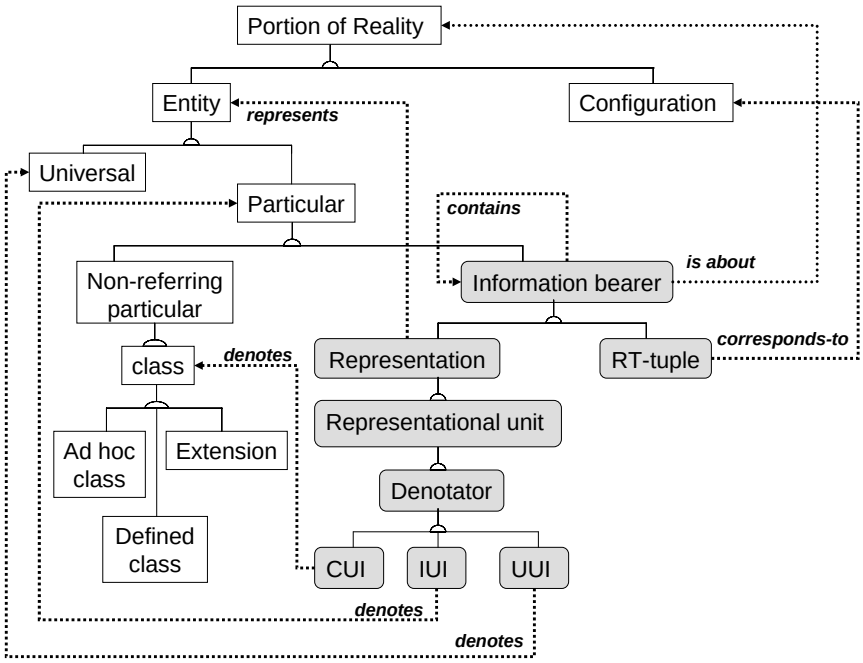

Figure 5: a referent tracking view on reality (all boxes) and the various entities that can be used to represent reality (gray, rounded boxes)

5.3 Principles for evolutionary ontology development (EOD)

The third item on the realist agenda in ontology development is the requirement that the structure of an ontology should mimic the structure of the PoR that is covered by the ontology. Granular Partition Theory (GPT) provides a formal account of what it means for a structure to mimic (or not) another structure [55]. GPT allows for instance an ontology that represents whales as fish to be recognized as incorrect, where an ontology that classifies whales as animals but not as mammals, while not incorrect, still to be what GPT calls '*locally non-transparent*'. GPT does however not provide a means

to quantify such differences, nor does it deal with issues such as whether it matters, for the purposes for which the ontology has been designed, whether whales are mammals, or what the reasons are for given sorts of mismatch. This is especially relevant in domains where our scientific understanding of reality is advancing rapidly and so that ontologies seeking to keep pace with these advances need to be updated frequently.

In [34], building further upon GPT, a metric is proposed to quantify the quality of ontologies on the basis of four dimensions: (1) type of structural mismatch as defined by GPT, (2) relevance for the purposes for which the ontology is designed, and whether structural mismatches arise (3) from a wrong or incomplete scientific understanding of the relevant parts of reality, or (4) from editorial mistakes.

5.3.1. Quantification of structural mismatches regarding representational units

As shown in **Table 7**, the current version of EOD is based on 17 possible configurations of match or mismatch – 2 more than in the original proposal [34] – which are divided into two groups, labelled '*P*' and '*A*', denoting respectively the presence or absence of an RU. Each group can further be subdivided into two smaller groups on the basis of whether the presence or absence of an RU in an ontology is justified ('*P*+' and '*A*+') or unjustified ('*P*-' and '*A*-').

The configurations reflect the different kinds of mismatch between what the ontology authors *believe* to exist or to be relevant, on the one hand, and matters of *objective* existence and *objective* relevance-to-purpose on the other. The encoding of a belief can be either correct (R+) or incorrect, either (a) because the encoding does not refer (¬R) or (b) because it does refer, but to a PoR other than the one which was intended (R-). The two configurations not considered in the original proposal [34] both involve an RU that denotes an intended and objectively existing PoR that, however, is already denoted by another RU in the ontology (R++).

Table 7: Typology of expressions included in and excluded from an ontology in light of relevance and relation to external reality

Configuration	Reality		Representation				Magnitude of error
			Authors' Belief		Encoding		
(1)	Objective Existence (2)	Objective Relevance (3)	In existence (4)	In relevance (5)	Intended encoding (6)	Type of reference (7)	(8)
P+1	Y	Y	Y	Y	Y	R+	0
A+1	N	-	N	-	-	-	0
A+2	Y	N	Y	N	-	-	0
P-1	N	-	Y	Y	Y	¬R	3
P-2	N	-	Y	Y	N	¬R	4
P-3	N	-	Y	Y	N	R-	5
P-4	Y	Y	Y	Y	N	¬R	1
P-5	Y	Y	Y	Y	N	R-	2
P-6	Y	N	Y	Y	Y	R+	1
P-7	Y	N	Y	Y	N	¬R	2
P-8	Y	N	Y	Y	N	R-	3
P-9	Y	Y	Y	Y	Y	R++	1
P-10	Y	N	Y	Y	Y	R++	2
A-1	Y	Y	Y	N	-	-	1
A-2	Y	Y	N	-	-	-	1
A-3	N	-	Y	N	-	-	1
A-4	Y	N	N	-	-	-	1

As an example, configuration P-1 would hold for an RU stating that 'whales are fish': the putative PoR does not exist – hence the 'N' in column (2) of **Table 7** – and therefore objective relevance does not apply, as indicated by the '-' in column (3). The authors of the ontology do however *believe* that

whales are fish and consider it to be relevant; therefore this configuration is marked by the presence of 'Y' in both columns (4) and (5). Finally, they use the representational machinery offered by the ontology correctly such that the RU is the intended representation – note the 'Y' in column (6) – but this in absence of a corresponding PoR, as indicated by '¬R' in column (7).

Of the 17 configurations, only 3 are desirable: P+1, which consists in the justified presence of an RU that correctly refers to a relevant PoR; and A+1 and A+2, which consist in the justified exclusion of an RU, either because there is no PoR to be referred to, or because this PoR is not relevant to the ontology's purpose. A-3 and A-4 are borderline cases, in which errors made by ontology authors are without deleterious effect, either because something that is erroneously assumed to exist is deemed irrelevant, or because something that is truly irrelevant is overlooked.

There are eleven different kinds of 'P' configurations of which, interestingly, only P+1 and P-6 refer correctly to a corresponding PoR: the former reflects our ideal case for presences; the latter is marred by the incorrect inclusion of an RU which lacks relevance. P-9 and P-10 also denote an existing and intended PoR, but the mistake here is that the ontology authors are not aware of their departure from the principle that for each entity in first-order reality there should be maximally one RU of a specific form.

The last column of **Table 7** shows the magnitude of the error committed when an RU reflecting a given type of configuration is included in or left out of an ontology as measured against its corresponding ideal configuration. Because these ideal configurations are P+1, A+1, and A+2, and because for any other configuration the '*corresponding*' ideal configuration is the one which has the same values in columns (2) and (3), the number of mistakes committed in P-4, P-5, P-9, A-1 and A-2 need to be measured against P+1. Similarly A+1 is the ideal configuration for P-1, P-2, P-3 and A-3, and A+2 for all the others. The magnitude of an error is calculated by counting the number of differences that a specific configuration exhibits with respect to its ideal configuration in each of the columns (4) to (7) of **Table 7**, with the additional rule that a non-intended encoding which denotes an existing and thus non-intended PoR – the presence of 'R-' in column (7) – counts double. This is because we judge that users of a terminology will be less likely to use RUs which denote nothing than RUs that denote non-intended PoRs: probably far more users will notice that an RU of the type 'whales are leprechauns' is a mistake – and thus never use that RU in some annotation – than there would be users that would notice the mistake in an RU of the type 'whales are fish'.

5.3.2. Quantification of structural mismatches regarding whole ontologies

Theoretically, it would now be an easy exercise to assess the quality of an ontology as a whole: we would have to (1) inspect each RU in the ontology to determine what match/mismatch configuration it exhibits, and (2) examine its coverage domain to see what relevant RUs are missing. Because the magnitude of a mistake in an undesirable configuration is maximally 5, we would give each best case configuration encountered a score of 5, while each deviation there from would receive the difference between 5 and the corresponding penalty for the corresponding sort of deviant case. The total score would be the ratio of the sum of the scores obtained for each present RU, over the sum of five times the number of RUs present and 4 times the number of RUs missing. The latter is because all missing RUs have an error magnitude of 1, and 5-1=4. The general formula is:

$$\frac{\sum_{i=1}^{n}(5-e_i)}{5n+4m} \qquad (1)$$

in which e_i stands for the magnitude of the error (if any) for a given corresponding *RU*, *n* for the number of *RUs* present in the ontology and *m* for the number of RUs unjustifiably absent. Note that in this study we did not assign a higher or lower error magnitude to unjustified absences that occur at the level of leaf nodes in a ontology as compared to absences at higher levels in the hierarchy.

The score itself can be viewed as a variation to the well-known recall and precision metric, but combined in but one metric and adjusted for the magnitude of the errors committed.

Table 8 gives an example of how this metric should be applied. Imagine three ontologies that provide a vocabulary for describing whales. All three ontologies have RUs for WHALE, FISH, ANIMAL and MAMMAL, but they differ in whether whales are asserted to be (1) fish (Ontology 1 - T1), (2) animals without further specification (Ontology 2 - T2), or (3) mammals (Ontology 3 - T3). In reality, of course, whales are mammals. We further assume, for the sake of the example, that the ontology authors did not make encoding mistakes: if there is a mistake in the ontology, then it is because their scientific understanding of reality is erroneous, not because they encoded a known fact erroneously. We also assume that all PoRs in the domain are relevant to the purposes for which the ontologies are built. When we then compare the three ontologies against the benchmark of reality, the latter being expressed in column (2) of **Table 8**, we see that T1 has one erroneous RU, which is an example of a mistake of type P-1, and one unjustified absence of type A-2; T2 exhibits the same unjustified absence, but in contrast to T1 it does not include an erroneous RU; T3, finally, mimics the structure of reality completely. For each RU in each ontology, the corresponding error magnitudes, if any, are shown in columns (4), (6) and (8). Applying the formula described above, this gives a quality score for T1 of 0.84, for T2 of 0.90 and for T3 of 1.00.

Table 8: Scoring the quality of ontologies using reality as benchmark

	Reality	Ontology 1		Ontology 2		Ontology 3	
RU(1)	Config. (2)	Config. (3)	Error (4)	Config. (5)	Error (6)	Config. (7)	Error (8)
animal	P+1	P+1	0	P+1	0	P+1	0
fish	P+1	P+1	0	P+1	0	P+1	0
whale	P+1	P+1	0	P+1	0	P+1	0
mammal	P+1	P+1	0	P+1	0	P+1	0
fish are animals	P+1	P+1	0	P+1	0	P+1	0
mammals are animals	P+1	P+1	0	P+1	0	P+1	0
whales are fish	A+1	P-1	3	A+1	0	A+1	0
whales are animals	P+1	P+1	0	P+1	0	P+1	0
whales are mammals	P+1	A-2	1	A-2	1	P+1	0
SCORE	8*5/ ((8*5)+(0*4)) = 1.00	((7*5)+(1*2))/ ((8*5)+(1*4)) =0.84		7*5/ ((7*5)+(1*4)) =0.90		8*5/ ((8*5)+(0*4)) =1.00	

Note that we took the justified absence of type A+1 (whales are fish) into account *only* because there is an RU (in T1) that posits the opposite. It is of course *not* a presupposition of our proposal that one should include all putative RUs which do not denote a corresponding PoR – e.g. that animals are fish, that animals are whales, that fish are mammals, that unicorns are leprechauns, and so forth – in any such assessment. Importantly, not doing so does not affect the magnitude of the overall score.

This can be seen in relation to T2 and T3 whose quality scores are not influenced by the fact that they do not contain an erroneous RU to the effect that whales are fish. This is one of the desirable mathematical properties that this metric exhibits, of which the complete characterization, however, is not yet completed.

Note also that this procedure reflects what might initially appear to be an unacceptable idealization, because determining the type of configuration an (included or excluded) RU is involved in depends upon two factors – objective relevance-to-purpose, and relation to objective reality – whose assessment is something which could be correctly carried out only by someone able to adopt the perspective of a god-like observer. Less idealistically, this god-like observer might be replaced by another ontology that is used as gold standard [56], and we adopt here a generalization of this latter approach by using successive versions of an ontology as the gold standard relative to its predecessors. This is motivated, as described further in detail, by the assumption that new versions of an ontology are better than previous ones, despite the possibility that with each version new errors are introduced. But if ontology curators take their work seriously, such errors are likely to be corrected in later versions, for instance on the basis of remarks from the community when the version is used in practice. It seems obvious that using other ontologies as gold standard has at least the same risk. Furthermore, if one is sure about the correctness of another ontology covering the same domain, why should one then bother to develop a new one?

5.3.3. Quality assessment of ontologies over successive versions

The minimal requirement for releasing an ontology as expressed in terms of the realist paradigm (though independent of whether or not authors of a given ontology endorse a realist view) is that its authors should assume in good faith that all its constituent expressions are of the P+1 type (requirement R1). A stronger requirement would be that the authors advance the ontology as complete, i.e. as containing RUs designating *all* PoRs deemed relevant to its purpose (requirement R2). Successive versions of an ontology should approximate ever more closely to this latter ideal. To exploit the paradigm completely, one could even argue that it should be part of the standard ontology authoring process to document any changes made in successive versions by means of the typology described in Table 7 [34]. This requires ontology authors to register whether or not the changes they introduced in a new version of the ontology are dictated by changes in (1) the underlying reality (requirement R3), (2) objective relevance of an included expression to the purposes of the ontology (requirement R4), (3) the ontology authors' understanding of each of these (requirement R5), and also by (4) the correction of encoding errors (requirement R6).

To see how the heuristic of using a new version of an ontology functions as surrogate for a god-like observer in relation to its predecessors, consider again the whale/fish example of **Table** 8. This time, however, we will consider T1, T2 and T3 to be versions of the same ontology, T3 being newer than T2, and T2 being newer than T1. The results of this interpretation are summarized in **Table 9**; with **Table 10** showing how the individual quality scores are calculated.

When the first version of the ontology (T1) is released, the authors assume in good faith that their work is correct, i.e. that all RUs denote the desired PoRs, and that all and only relevant RUs are present. They might believe that some RUs are missing, but of course, they have no clue which ones, otherwise they would have been included. Therefore, version T1 at time t_1 was assumed to be 'state of the art' and therefore of quality 1.00, the maximal attainable score. At time t_2, however, the authors discover that whales are not fish and they make the corresponding RU 'obsolete'. Note that making an RU obsolete by giving the reason for the change, is preferable to just removing it: if, indeed, the only change introduced between T2 and T1 would be the deletion of the RU that whales are fish, external

auditors might wonder whether (1) the deletion is an omission brought about by an encoding error, in which case the RU which was believed to be of type P+1 at t_1 has to be believed to be of type P-2 at t_2, or (2) a deletion based on a conscious decision either (2a) that whales are still to be considered to be fish, but that the RU is not relevant for the purposes for which the ontology is being built, hence consisting in an A-3 type of mistake, or (2b) that the right sort of discovery was made and thus the original RU was of type P-1. Because the latter is the case, the quality score of T1 at t_1 can be recalculated according to the state of the art reached at t_2 using Eq. (1).

A similar analysis can be carried out at t_3, but now applied to both T1 and T2; in general, each new version of a terminology allows us to assess the quality of all previous versions of the ontology in light of the state of the art reached when the new version is released (see Table 11).

Table 9: Scoring the quality of ontologies using new versions

	Time t1		Time t2				Time t3					
	T1		T1		T2		T1		T2		T3	
	C.	E.	C.	E.	C.	E.	C.	E.	C.	E.	C.	E.
animal	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
fish	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
whale	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
mammal	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
fish are animals	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
mammals are animals	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
whales are fish	P+1	0	P-1	3	A+1	0	P-1	3	A+1	0	A+1	0
whales are animals	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0	P+1	0
whales are mammals	-	-	-	-	-	-	A-2	1	A-2	1	P+1	0
SCORE	1.00		0.93		1.00		0.84		0.90		1.00	

Table 10: Calculation of quality scores for ontology versions at different times

Ontology	Time of assessment	Formula for quality score	Quality Score
T1	t1	(8*5)/(8*5)	1.00
	t2	((7*5)+(1*2))/(8*5)	0.93
	t3	((7*5)+(1*2))/((8*5)+(1*4))	0.84
T2	t2	(7*5)/(7*5)	1.00
	t3	(7*5)/((7*5)+(1*4))	0.90
T3	t3	(8*5)/((8*5)+(0*4))	1.00

Table 11: Views on the quality of a ontology through successive versions

	Time		
Ontology version	t1	t2	t3
T1	1.00	0.93	0.84
T2	-	1.00	0.90
T3	-	-	1.00

6 Disease Ontology

Effective knowledge representation requires the use of standardized nomenclatures to ensure both shared understanding between people and interoperability between computer systems. Unfortunately, many existing biomedical vocabulary standards do not consistently distinguish between *disease*, clinical and pre-clinical *manifestations of disease*, and *diagnosis* [57]. In [58], a logical framework for describing the initiation and progression of disease and the associated evaluation procedures and information artifacts based on BFO and Referent Tracking is outlined. A view of disease is defended involving in every case some physical basis within the organism that bears a disposition toward malfunctioning and thus towards the initiation and execution of pathological processes.

6.1 Motivation

A framework is needed that encompasses representations of diseases, the manifestations of the underlying disorders, and the ways these manifestations are recognized and interpreted in the clinic. Inspection reveals that these issues thus far have not been adequately dealt with in standard terminology and ontology resources. The National Cancer Institute Thesaurus, for example, identifies 'Chronic Phase of Disease' as a subtype of 'Finding', which it defines as: '*Objective evidence of disease perceptible to the examining physician (sign) and subjective evidence of disease perceived by the patient (symptom)*' (http://nciterms.nci.nih.gov). This definition implies, however, that the chronic phase of a disease does not exist except as one or other form of evidence, and thus as perceived in some way. This example illustrates a common conflation between processes on the side of the organism on the one hand, and the observation of such processes on the other. This conflation is for a number of reasons problematic, as is revealed when, in the context of adverse events, more precisely the detection and prevention thereof, we need to investigate the relations between clinical phenomena and underlying causes.

6.2 Defined and Undefined Terms

While it is generally good practice to provide precise definitions for the terms assembled in an ontology, in any logically coherent approach to definition, some terms must remain undefined in order to avoid circularity or infinite regress. The undefined terms used in what follows are of three sorts: either

1. they are non-technical terms derived from ordinary English; or
2. they are technical terms derived from basic science (for example, 'organism', 'molecule'); or
3. they are primitive terms specific to the domain of interest.

Terms in group (3) – specifically: 'physical basis', 'bodily feature', 'clinically abnormal', 'homeostasis', 'disposition', and 'realization' – require special attention as they provide an indispensable foundation for the definitions provided. While, ex hypothesi, we cannot provide definitions for these terms, we can provide some elucidation and illustrative examples.

Below are the representational units relevant for the ReMINE project. We indicate for each unit the level of reality concerned and the sort of generic portion of reality it denotes.

***Physical basis of health and disease* (L1, Defined Class):** this includes any collection of the physical components in the interior and of the organism, whether healthy or diseased, at any level of granularity from a single nucleotide to an arm. Thus it includes a liver cell that has been infected by a virus, a polyp, a clot, a displaced disk.

Bodily feature **(L1, member of Composite Class):** any member of the composite defined class comprising such collections of not only physical components but also associated bodily qualities and processes in an organism, both physiological and pathological.

Clinically abnormal bodily feature **(L1, Bodily feature):** any bodily feature of or in an organism that is

- causally linked to an elevated risk either of pain or other feelings of illness, or of death or dysfunction),
- not part of the life plan for an organism of the relevant type (unlike pregnancy or menopause), and
- such that the elevated risk exceeds a certain threshold level.

The term 'clinical' in the above has the following overlapping meanings:

(1) clinical as a temporal attribute of a disease process (meaning, roughly, late, overt, a matter of signs and symptoms and thus of what is observable by patient and clinician); and

(2) clinical as an attribute of an examination (meaning, roughly, relating to what doctors do with patients in their clinics).

It is worth noting that this treatment of 'abnormal' is distinct from those statistical treatments which do not take account of the overlap in the distribution of test results between normal and abnormal populations or of normal distribution extremes. This treatment takes also 'normal variants' into account since they satisfy the 2nd criterion, but not the 1st and 3rd.

Homeostasis **(L1, disposition):** disposition in an organism to regulate its bodily processes in such a way as (1) to maintain bodily qualities within a certain range or profile and (2) to respond successfully to departures from this range caused by internal or external influences.

When bodily processes yield qualities outside the homeostatic range, then the organism initiates processes designed to return the qualities to a value within this range. In some cases, homeostasis can be lost and then re-gained at a level that is clinically abnormal, for example in the case of adaptation to major injury. In other cases the organism will pass a point where it falls irreversibly outside the realm of homeostasis.

'***Disposition***' is a technical term taken from BFO which in the context here denotes an entity on the side of an organism in virtue of which it will initiate certain specific sorts of processes whenever certain associated sorts of preconditions are satisfied. Examples are: our disposition to crave liquid following dehydration; the disposition of an epithelial cell in the G2 phase of the cell cycle to become diploid following mitosis. In any organism there is a wide variety of dispositions, some associated with health, others with disease and other sorts of clinical abnormality. The term '***realization*****'** denotes the process through which a disposition is realized. Each disposition in an organism has a physical basis in bodily structures of an associated type.

In what follows, we pursue a view of disease as resting in every case on some (perhaps as yet unknown) physical basis. When, for example, there is a persistent elevated level of glucose in the blood of a patient, this is because (1) some physical structure or substance in the patient is disordered (loss of beta cells in pancreatic islets) as a result of which (2) there exists a disposition (diabetes) for the organism to act in a certain abnormal way. The disposition in question is realized by the initiation and execution of specific pathological processes (diabetic nephropathy) whose manifestations can be recognized as signs of the disorder (proteinuria).

Thus as a result of an etiological process, a physical change occurs in the healthy individual giving rise to a disorder, which initiates the development of a series of manifestations that are initially

undetectable without the use of special instruments (pre-clinical) and then become detectable as symptoms and signs. These clinical manifestations of disease constitute in their totality the clinical phenotype for the given disease as instantiated in this specific patient. They can be observed through physical examination and laboratory testing of specimens derived from the patient, the results of which can be recorded in the medical record as the clinical picture. The clinical picture is interpreted by the physician in arriving at a diagnosis, which serves in turn as the foundation for the development of a patient management plan.

Some advantages of this account of disease as dispositions rooted in physically disordered structures in the organism which are realized in pathological processes are that, in contrast to definitions of disease in terms of signs and symptoms or in terms of disease processes, it helps us to do justice (1) to the existence of pre-clinical manifestations of disease, (2) to the different combinations of disease and predispositions to disease in the form of elevated risk, and (3) to the fact that the disease course and the clinical picture may vary widely between patients even where the patients in question have the same disease.

Thus, this framework should support more sophisticated approaches to the treatment of patient data for clinical decision support and specifically for managing risks against patient safety.

6.3 Entities and configurations on the side of the patient

Homeostatic Range **(L1, portion of reality) =def. –** The range of values for a set of bodily feature types whose maintenance is continuously sought by an organism in homeostasis (e.g. 65 – 110 mg glucose/dL serum).

The homeostatic range for a given organism will vary in light of environmental and behavioral changes, for example to reflect raised heart beat frequency while running.

Abnormal Homeostasis **(L1, homeostasis) =def. –** Homeostasis that is clinically abnormal for an organism of a given type and age in a given environment (e.g. maintenance of high blood pressure).

Normal Homeostasis **(L1, homeostasis) =def. –** Homeostasis that is not clinically abnormal.

Disorder **(L1, portion of reality) =def**. – A configuration or collection of physical components in the interior or surface of the organism, including cells and other physical structures as well as portions of bodily substance, that is clinically abnormal (e.g. a tumor, an infected cell, a prion molecule in the brain, mutated genomic DNA, endotoxin in blood, blood with reduced blood cortisol levels causing adrenal crisis).

It is disorders so defined here that are the physical basis of disease. Diseases come into existence because a bodily structure is malformed or because a bodily substance such as blood is affected by the presence of a pathogen or toxin or by the absence of blood antibodies, in ways that lead to impairment of normal functioning.

We can distinguish a less and a more inclusive usage of the term ‘disorder’. When the term is taken in the former sense, each single cell within a tumor may be counted as a disorder in its own right. When the term is taken in the latter sense, the disorder is the tumor as a whole, the maximal collection of all disordered cells. It is the latter usage that corresponds most closely to the language used by clinicians.

Pathological Process **(L1, process) =def. –** A biological process in an organism that is clinically abnormal (e.g. transient inflammation in response to bacterial infection).

We can distinguish among pathological processes between those that are changes in the way a normal physiological function is realized (for example hyperventilation) and those that have no normal physiological counterpart (for example acute inflammation).

***Disease* (L1, disposition) =def. –** A disposition (i) to undergo pathological processes that (ii) exists in an organism because of one or more disorders in that organism (e.g. epilepsy as a disease that disposes to the occurrence of seizures (pathological process) due to an underlying abnormality in the neuronal circuitry of the brain (physical basis); AIDS as a disease that disposes to opportunistic infections that take advantage of a weakened immune system).

A predisposition is a disposition to acquire a further disposition. Some diseases, for example AIDS, are predispositions to further diseases.

***Predisposition to Disease of Type X* (L1, disposition) =def. –** A disposition in an organism that constitutes an increased risk of the organism's subsequently developing the disease X.

***Etiological Process* (L1, process) =def. –** A process in an organism that leads to a subsequent disorder (e.g. toxic chemical exposure resulting in a mutation in the genomic DNA of a cell).

The etiological process creates the physical basis of the disposition to pathological processes which is the disease. Some diseases are such that a patient can suffer from what is qualitatively the same disease on two distinct occasions **–** for example two successive bouts of influenza. We then say that the patient has two distinct disease instances of the same disease type. These successive bouts are differentiated by their etiology in the sense that their respective physical bases are caused by distinct processes, each prior to the successive disease instances which they cause.

The above definition implies a distinction between true etiological determinants of a disease from those that modify the presentation and course of the disease, for example by serving as the causes of clinical phenotypes such as inflammation common to many diseases.[2]

***Disease Course* (L1, collection of processes) =def. –** The totality of all realizations of a given disease instance.

***Transient Disease Course* (L1, disease course) =def. –** A disease course (e.g. a bout of flu) that terminates in a return to normal homeostasis.

***Chronic Disease Course* (L1, disease course) =def. –** A disease course that (a) does not terminate in a return to normal homeostasis and (b) would, in the absence of intervention, fall within an abnormal homeostatic range (e.g. intermittent seizures in a person suffering from epilepsy, acquired deafness).

***Progressive Disease Course* (L1, disease course) =def. –** A disease course that (a) does not terminate in a return to homeostasis and (b) would, in the absence of intervention, involve an increasing deviation from homeostasis (e.g. malignant cancer).

***Infection* (L1, disorder) =def. –** A disorder of a type which involves changes brought about by a pathogenic organism within a host organism which lead to pathogen persistence and/or pathogen duplication.

***Infectious Disease* (L1, disease) =def. –** A disease caused by an infection.

Examples: *transient*: seasonal flu; *chronic*: genital herpes; *progressive*: Ebola-virus-mediated hemorrhagic fever.

[2] Further work is needed here: some etiological processes are also pathological processes: a disease may cause processes that lead to pain and dysfunction which themselves turn the patient depressive. Or overweight because of chronic hypothyroidism causes knee and hip arthrosis. Thus it seems that there is no 'essence' in pathological or etiological process, and that therefore they are to be seen as defined/composite classes.

Secondary Infection **(L1, disorder) =def. –** A disorder consisting in (1) the presence of a pathogenic organism within a host organism that (2) leads to infection in virtue of (3) a predisposition to disease that exists in virtue of (4) a prior infection with a pathogenic organism of a different kind (e.g. cryptosporidiosis in a patient suffering from AIDS)

6.4 Definitions of Terms Relating to Clinical Evaluations

In many cases, patients harbour disorders for many years before the associated dispositions are realized in changes which cross the threshold of observability. The latter are frequently first recognized by patients (symptoms) and subsequently observed by clinicians (signs). Although the terms 'sign' and 'symptom' are frequently used in this way to distinguish the source of information, since both represent manifestations (observed realizations) of disease, the distinction may be of limited utility. We believe that a more rigorous treatment of the distinction would be through the explicit representation of the agents involved in different sorts of observations (thus: direct clinician observation, patient report, report of some other observer, laboratory test, and so on). However, because the sign/symptom distinction is made routinely by clinicians in the conduct of patient care, we include definitions for 'sign' and 'symptom' which are conformant with the other definitions here provided.

Sign **(L1, bodily feature) =def. –** A bodily feature of a patient that is observed in a physical examination and is hypothesized by the clinician to be of clinical significance.

The sign provides evidence to believe in the existence of an underlying disorder. Neither signs nor symptoms form a natural kind, but are members of composite classes – fiat collections of bodily features delineated in reflection of certain cognitive practices on the parts of clinicians and patients.

A distinction which has perhaps to be made – this is an issue that needs further work – is that between 'sign' and 'sign of'. So an observed clinically abnormal bodily feature is for two clinical observers always a 'sign' but not necessarily a sign of the same disorder.

Vital sign **(L1, bodily feature)** =**def.** – A physical sign in which the presence of a non-zero value is standardly considered to be an indication that the organism is alive. The relative levels of vital sign values are often used as measures that can indicate the presence of disease.

Symptom **(L1, bodily feature) =def. –** A quality of a patient that is observed by the patient and is of the type that is hypothesized by the patient as a realization of a disease.

On some readings of the term, symptoms – paradigmatically pain and other feelings and sensations – are such that they can be observed *only* by the patient.

Clinical History Representation **(L3, information entity) =def. –** A series of statements about the clinical history of a person.

Clinical history **(L1, collection of processes):** collection of health-relevant features of a patient and of a patient's family, including, when present, a collection of a patient's disease courses.

Also a patient that never went to the doctor has a clinical history, but probably not a clinical history representation.

Clinical History Taking **(L1, process) =def.** – An interview in which a clinician elicits a clinical history from a patient or from a third party who is authorized to make health care decisions on behalf of the patient.

Physical Examination **(L1, process) =def. –** A sequence of acts of observing and measuring bodily features of a patient performed by a clinician.

Measurements may occur with and without elicitation.

Laboratory Test **(L1, process) =def. –** A measurement assay that has as input a specimen derived from a patient, and as output a result that represents a quality of the specimen.

Laboratory Finding **(L3, information entity) =def. –** A representation of a quality of a specimen that is the output of a laboratory test and that supports an inference to an assertion about some quality of the patient.

Normal Value **(L3, information entity) =def. –** A range of values for a given quality reported in a lab report and asserted by the testing lab or the kit manufacturer to be normal on the basis of a statistical treatment of values from a reference population.

Clinical Finding **(L3, information entity) =def. –** A representation of a bodily feature of a patient that is the output of a clinical history taking, of a physical examination, of some other investigation, or some combination thereof.

Preclinical Finding **(L3, information entity) =def. –** A representation of a bodily feature of a patient that is (1) recorded by a clinician because the feature is hypothesized to be of clinical significance and (2) refers to features obtaining in the patient prior to their becoming detectable in a clinical history taking or physical examination.

Manifestation of a Disease **(L1, bodily feature) =def. –** A bodily feature of a patient that is (a) a deviation from clinical normality that exists in virtue of the realization of a disease and (b) is observable (including observable through elicitation of response and through the use of special instruments).

Preclinical Manifestation of a Disease **(L1, bodily feature) =def. –** A manifestation of a disease that exists prior to its becoming detectable in a clinical history taking or physical examination.

Clinical Manifestation of a Disease **(L1, bodily feature) =def. –** A manifestation of a disease that is detectable in a clinical history taking or physical examination.

Phenotype **(L1, bodily feature or collection thereof) =def. –** A bodily feature or constellation of bodily features of an organism determined by the interaction of the genetic make-up and environment of the organism.

Clinical Phenotype **(L1, phenotype) =def.** – A clinically abnormal phenotype.

Disease Phenotype **(L1, phenotype) =def. –** A clinically abnormal phenotype that is characteristic of a single disease.

Note that, according to this definition, a disease phenotype can exist without being observed. Indeed, as technology advances, our ability to detect the underlying components of a disease phenotype expands. What we might think of as the full disease phenotype would incorporate the abnormal phenotypes realized at each stage in the development of the disease.

As with ‘disorder’, so also with ‘phenotype’ we can distinguish a less and a more inclusive reading. Under the former a disease phenotype may be a single quality, or a single abnormality type; under the latter it may be the maximal constellation of all the phenotypes associated with a given disease.

Clinical Picture **(L3, information entity) =def.** – A representation of a clinical phenotype that is inferred from the constellation of laboratory, image and clinical findings available to the clinician about a given patient.

Diagnosis **(L2, belief) =def. –** A conclusion of an interpretive process that has as input a clinical picture of a given patient and as output an assertion (diagnostic statement) to the effect that the patient has a disease of such and such a type.

The diagnostic process is typically iterative: the clinician is forming hypotheses during history taking, testing them during additional history taking, forming new hypotheses as a result, testing these during physical exam, forming new hypotheses as a result, and so on.

7 An Ontological Analysis of 'Adverse Event'

7.1 Objective, design and methods

Our goal was to bring clarity to the terminological wilderness that grew out of all the efforts documented in [11]. 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* […]" [59] (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.

7.2 Results

7.2.1. Core representational units

Table 12 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.

7.2.2. 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 [11]. As an example, **Table 13** 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.

Table 12: **Universals and Defined Classes for the adverse events domain.**

	Denotation	Class Type	Particular Type	Description (role in adverse event scenario)
			Level 1	
C1	subject of care	DC	independent continuant	person to whom ***harm*** might have been done through an ***act under scrutiny***
C2	act under scrutiny	DC	act of care	***act of care*** that might have caused ***harm*** to the ***subject of care***
C3	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***
C4	care giver	DC	independent continuant	person that performed an ***act of care*** directed to the ***subject of care***
C5	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***
C6	involved structure	DC	independent continuant	anatomical structure (of the ***subject of care***) involved in an ***act of care***
C7	structure change	U	process	change in an anatomical structure of a person
C8	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
C9	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
C10	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
C11	care effect	DC	integrity change	***integrity change*** brought about by an ***act of care***
C12	subject investigation	DC	process	looking for a ***structure change*** in the ***subject of care***
C13	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***
C14	care intention	DC	dependent continuant	intention of a ***care giver*** that motivates him towards an ***act of care***
			Level 2	
C15	observation	DC	dependent continuant	cognitive representation of a ***structure change*** resulting from an act of perception within a ***subject investigation***
C16	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***
C17	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	
C18	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***

The relationships employed in composing representations of properties in this Table are drawn from [38, 49]. We preserve the formatting conventions proposed in [38], 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.

We also allow for temporal annotations additional to those described in [38], at the same time remaining faithful to EN 12388:2005: Health Informatics – Time Standards for Healthcare Specific Problems [60].

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 [13], it would be either ***#12*** or ***#13***. The counterintuitive nature of the latter case has its roots in certain conflations in the HL7 RIM [61], 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.

Table 13**: 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}

7.3 Discussion

Already a very superficial analysis of the definitions in **Table 2** (page 2) 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*'.

7.3.1. 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 expected 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.

7.3.2. 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 [62]), 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 [19], 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 [63, 64], it is considered bad practice not to do such a test at regular intervals [65]. 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.

7.4 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.

8 Domain Ontology for Managing Risks against Patient Safety (work in progress)

8.1 General context

Under the view of the ReMINE model to the effect that an adverse event is an 'incident [that] occurred during the past and [is] documented in a database of adverse events' ([16], see also section 3.2 p2), and the objective of the project not only to use this database to prevent adverse events but also to advance the state of the art in adverse event management, we designed the ontology to cover (1) the first-order reality in which incidents happen to patients, and (2) the corresponding third-order reality, i.e. the information artifacts (L3)[3] which are intended to describe the first-order reality.

First-order entities represented in the RAPS Domain ontology are primarily discussed in chapters 6 and 7. We provide here more background on the representations.

For this ontology to be applicable, such a database – it may be empty – has to exist in the environment in which the ontology is used. This database must conceptually consist of at least two parts: an adverse event repository (L3) that contains denotators (L3) that stand proxy for the incidents in reality that are considered to qualify as adverse events, and a database of relational expressions (L3) that describe to what other entities these incidents are related. We will use the name adverse event knowledge base (L3) for the latter.

Indeed, evidence from external sources combined with experiences in the ReMINE pilot sites about the occurrence and prevention of risks against patient safety lead to the creation of various sorts of information artifacts (L3) [66] some of which – the ones that concern us the most here and which we call executable information artifacts (L3) – are used to steer guideline execution and protocol monitoring applications.

Faithful information artifacts, i.e. information artifacts that are faithful to reality, contain in one or other form representations (L3) about portions of reality, the latter being either portions of first-order reality (L1), beliefs (L2), or other information artifacts. The smallest parts of faithful information artifacts are denotators which stand proxy for corresponding entities in reality (either L1, L2 or L3 entities). Denotators of faithful information artifacts are usually combined in relational expressions (L3) that correspond to configurations.

An executable information artifact is created by an executable information artifact author (L1) and is intended to be a pure syntactic transformation from some other information artifact or combination of information artifacts with the goal to be used in an algorithm for automatic reasoning. If this creation has been done error free, and if the information artifact(s) that was (were) used as source were faithful to reality, then the resulting executable information artifact(s) represent(s) the same portion of reality as the one referred to by the source information artifact(s).

Of course, things may go wrong, not only during the transformation process, but also in relation to the faithfulness of the information artifacts that were used as source. For some information artifacts it might not be known whether they are faithful to reality, but we assume that if an information artifact is qualified as an active information artifact (L3) in the adverse event knowledge base, it is assumed to be faithful by at least one member of the RAPS system authority (L1) and at least one RAPS system user (L1). If no such member or user endorses the information artifact anymore, it becomes qualified as retired information artifact (L3). Keeping track at all times of which RAPS system authority members and RAPS system users endorse what information artifacts during what periods of time, as well as the history of activation and retirement of information artifacts holds great promises for quality assurance purposes [34, 67, 68].

[3] L1, L2 and L3 are used to indicate the level of reality involved (see section 5.1, p2)

8.2 RAPS Domain ontology definitions

In this section we provide semi-formal descriptions of what the representational units in the RAPS Domain Ontology stand for. Following the OBO-Foundry's principle of modularity, we only provide representational units for those entities that are not yet described in officially released versions of other ontologies to which we link.

We present the representational units in three tables, one for each level of reality. For each unit, we provide:

- a term ('Denotation'),
- the type of unit involved, i.e. 'U' for universal, 'DC' for Defined Class, 'CC' for Composite Class ('Unit type'),
- what universals the particulars instantiate ('Particular Type'),
- a semi-formal description, and optionally,
- comments.

We use in the descriptions and comments the following font and type setting conventions:

- red font: level-1 entities defined in the RAPS Domain Ontology,
- green font: level-2 entities defined in the RAPS Domain Ontology,
- blue font: level-3 entities defined in the RAPS Domain Ontology,
- magenta font: names (or syntactic variation thereof in order to write well-formed grammatical sentences) for formal relations,
- **bold**: entities defined in ontologies we link to.

The tables are sorted alphabetically on the denotator column such that the descriptions for units used in descriptions and comments can be found easily.

8.2.1. Level-1 entities

Denotation	Unit Type	Particular Type	Description (role in risk management)	Comment
act of care	U	process	**process** (1) which has agent a care giver and (2) underwent by a subject of care, and (3) is motivated by an underlying disease and a care intention	• (3) excludes that **processes** whose agents are care givers but that are not performed under the care giver role would be qualified as acts of care (e.g. a doctor hurting a patient in a car accident on the parking lot of a care facility)
act under scrutiny	DC	act of care	act of care which is member of process under scrutiny	
adverse event	DC	process	**process** denoted by a denotator in a RAPS adverse event repository	• If the RAPS adverse event repository were faithful to reality, each member of adverse event would be a member of harm. • If he RAPS adverse event repository were *locally complete*, each member of harm that occurred in the RAPS system's realm in which the RAPS adverse event repository is installed would be a member of adverse event
anatomical integrity change investigation	DC	process	**process** undergone by a **human being** which has agent a care giver to detect a change in anatomical structure in that **human being** resulting in an anatomical structure change observation that inheres in the care giver	• to avoid confusions of the sort found in HL7-RIM [61], we distinguish explicitly between the act of observing, what is observed and what is cognitively invoked in the observer through the act of observing

Denotation	Unit Type	Particular Type	Description (role in risk management)	Comment
anatomical integrity change	U	change in anatomical structure	change in anatomical structure causing a change in its anatomical structure integrity	
anatomical structure integrity	U	quality	**quality** of an **anatomical structure** determining the range of circumstances of the sort typically occurring in the life of the **human being** under which the **anatomical structure** of that **human being** would become dysfunctional or cause dysfunction in another **anatomical structure**	• **quality** is defined in BFO • **human being** and **anatomical structure** are defined in the Foundational Model of Anatomy (FMA) [39] • in each **anatomical structure** inheres an instance of anatomical structure integrity
availability	U	ability	ability to participate in a process	• **ability** is defined in BFO ([62])
bodily harm	U	anatomical integrity change	an anatomical integrity change bringing about a deterioration of an **anatomical structure**'s anatomical structure integrity	• **human being** and **anatomical structure** are defined in the Foundational Model of Anatomy (FMA) [39] • each instance of bodily harm is a member of harm
bodily harm assessment	U	process	**process** which has agent a care giver to determine whether an anatomical structure change observation is faithful to reality, and if so, whether the change in anatomical structure which is denoted by the anatomical structure change observation is a bodily harm	
harm assessment	U	process	**process** which has agent a person to determine whether another **process** is a harm	
care effect	DC	anatomical integrity change	anatomical integrity change caused by an act of care	
care giver	DC	human being	**human being** in which inheres an instance of care giver role	
care giver role	U	role	**role** inhering in a **human being** mandated to be the agent of acts of care	
care intention	DC	dependent continuant	intention inhering in a care giver that motivates him towards an act of care	
change in anatomical structure	U	process	change in an **anatomical structure** of a **human being**	• **human being** and **anatomical structure** are defined in the Foundational Model of Anatomy (FMA) [39]
harm	DC	process	a **process** resulting in an expansion in the range of circumstances of the sort occurring in the history of a continuant under which that **continuant** would participate in **processes** involving some sort of loss or detriment, whether physically, functionally, socially, economically, etc.	• **process** and **continuant** are defined in BFO • each instance of bodily harm is a member of harm
harmless process	CC	process	**process** that does not cause harm	• required to explain the REMINE term 'no harm'
hazard	U	potentiality	**potentiality** to cause harm	• **potentiality** is defined in BFO ([62])
healthcare associated infection	DC	infection	**infection** in a **human being** who participated in a **process** that (a) occurred in a setting delivering health services or (b) is an act of care	• **infection** is (being) defined in the Disease Ontology (see chapter 6 and [58])
history of a continuant	U	processual entity	a collection of **processes** in which a **continuant** participates or has participated	• **process** and **continuant** are defined in BFO • each **continuant** has a history of a continuant in which it participates
involved structure	DC	anatomical structure	**anatomical structure** which is part of the subject of care which undergoes an act of care	• the involved structure may or may not participate in the act of care (an example of the latter would be the skin undergoing an allergic reaction after the subject of care ingested a drug)

Denotation	Unit Type	Particular Type	Description (role in risk management)	Comment
mild bodily harm	U	bodily harm	bodily harm bringing about a mild expansion in the range of circumstances of the sort typically occurring in the life of the **human being** under which the **anatomical structure** of that **human being** would become dysfunctional or cause dysfunction in another **anatomical structure**	
moderate bodily harm	U	bodily harm	bodily harm bringing about a moderate expansion in the range of circumstances of the sort typically occurring in the life of the **human being** under which the **anatomical structure** of that **human being** would become dysfunctional or cause dysfunction in another **anatomical structure**	
person age	U	quality	quality of a **human being** determined by the amount of time elapsed since the **human being**'s birth	
process under scrutiny	DC	process	**process** denoted by a denotator which participates in computational **processes** to assess whether that **process** is a member of adverse event.	
severe bodily harm	U	bodily harm	bodily harm bringing about a large expansion in the range of circumstances of the sort typically occurring in the life of the **human being** under which the **anatomical structure** of that **human being** would become dysfunctional or cause dysfunction in another **anatomical structure**	
staff	DC	collection	collection of all staff members of a legal entity	
staff adequacy	DC	ability	ability of a staff member or team to deliver a service	
staff availability	DC	availability	availability of a staff member or team to deliver a service	
staff member	DC	human being	**human being** given the **role** to perform activities on behalf of a legal entity	• each staff member is member of staff
subject of care	DC	human being	**human being** undergoing acts of care	• may be qualified as '*patient*' although this term might not be appropriate in certain care facilities
subject under scrutiny	DC	human being	**human being** denoted by a denotator in which inheres the **disposition** to participate in computational **processes** to assess whether that **human being** participates in **processes** which are members of adverse event.	• Does not need to be the patient, but e.g. a visitor to or employee of the care facility using the RAPS system. • Does not need to be through an act of care, e.g. slipping on a wet floor
team	DC	collection	collection of staff members of a legal entity	• it is not the case that each staff member is member of a team
unavailable continuant	DC	continuant	**continuant** which lacks availability	• is applicable to devices, staff, drugs, operating rooms, beds, and so forth
underlying disease	DC	disease	the **disease** in the subject of care which is part of what serves to motivate performance of the act of care	• **disease** is (being) defined in the Disease Ontology (see chapter 6 and [58])

8.2.2. Level-2 entities

Denotation	Unit Type	Particular Type	Description (role in risk management)	Comment
anatomical structure change observation	DC	dependent continuant	cognitive representation of a change in anatomical structure resulting from an anatomical integrity change investigation	
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	
harm diagnosis	DC	dependent continuant	cognitive representation resulting from a harm assessment, and involving a conclusion that a certain process is or is not a harm	
positive harm diagnosis	DC	dependent continuant	harm diagnosis involving a conclusion that a certain process is a harm	
bodily harm diagnosis	DC	dependent continuant	cognitive representation, resulting from a bodily harm assessment, and involving a conclusion to the effect that a change in anatomical structure is or is not a bodily harm	
positive bodily harm diagnosis	DC	dependent continuant	harm diagnosis involving a conclusion that a change in anatomical structure is a bodily harm	

8.2.3. Level-3 entities

Denotation	Unit Type	Particular Type	Description (role in risk management)	Comment
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 in which participate a subject of care	
denotator	U	representational unit	**representational unit** denoting an **entity**	
executable information artifact	U	generically dependent continuant	information artifact which is a pure syntactic transformation from some other information artifact or combination of information artifacts, in which inheres the **disposition** to participate in computational **processes**	
information artifact	U	generically dependent continuant	**generically dependent continuant** which is intended to be about a **portion of reality**	• such artifact is generically dependent because it can exist in multiple copies (e.g. the 'same' pdf file on different hard drives)
RAPS adverse event repository	U	information artifact	collection formed by parts which are denotators that denote **processes** that started in relation to an act of care and are believed to be a member of harm	

9 Future work

The ontology presented thus far in chapter 8 is what is called the 'upper ontology' for the domain. It will expand considerably based upon the priorities that will be set forward in the project.

Future work will include:

- expressing the parts of the terminology presented in [16] that are relevant for the pilot site implementations in terms of the ontology, as well as any future updates of this terminology;
- mapping to the vocabulary used in the domain background materials described in section 3.3 and appendices 1 and 2;
- mapping to the databases with patient information used at the ReMINE pilot sites
- introducing adaptations in line with the evolutions in the OBO-Foundry ontologies.

10 Appendix 1: Definitions from the ICPS

Accountable	being held responsible
Actions taken to reduce risk	actions taken to reduce, manage or control the harm, or probability of **harm** associated with an **incident**.
Adverse event	an **incident** which results in **harm** to a patient.
Adverse reaction	unexpected harm resulting from a justified action where the correct process was followed for the context in which the event occurred.
Agent	a substance, object or system which acts to produce change.
Ameliorating action	an action taken or **circumstances** altered to make better or compensate any **harm** after an **incident**.
Attributes	qualities, properties or features of someone or something.
Circumstance	any factor connected with or influencing an **event**, **agent** or person(s).
Class	a group or set of like things.
Classification	an arrangement of **concepts** into **classes** and their subdivisions to express the **semantic relationships** between them.
Concept	a bearer or embodiment of meaning.
Contributing Factor	a **circumstance**, action or influence which is thought to have played a part in the origin or development of an **incident** or to increase the **risk** of an **incident**.
Degree of harm	the severity and duration of harm, and the treatment implications, that result from an **incident**.
Detection	an action or **circumstance** that results in the discovery of an **incident**.
Disability	any type of impairment of body structure or function, activity limitation and/or restriction of participation in society, associated with past or present **harm**.
Disease	a physiological or psychological dysfunction.
Error	failure to carry out a planned action as intended or application of an incorrect plan.
Event	something that happens to or involves a **patient**.
Harm	impairment of structure or function of the body and/or any deleterious effect arising there from.
Hazard	a **circumstance, agent** or action that can lead to or increase risk.
Health	a state of complete physical, mental and social wellbeing and not merely the absence of **disease** or infirmity.
Healthcare	services received by individuals or communities to promote, maintain, monitor or restore **health**.
Healthcare-associated harm	**harm** arising from or associated with plans or actions taken during the provision of healthcare rather than an underlying **disease** or **injury**.
Incident characteristics	selected **attributes** of an **incident**.
Incident type	a descriptive term for a category made up of incidents of a common nature grouped because of shared, agreed features.
Injury	damage to tissues caused by an **agent** or **circumstance**.

Mitigating factor	an action or **circumstance** which prevents or moderates the progression of an **incident** towards harming a **patient**.
Near Miss	an **incident** that did not cause **harm**.
Organizational Outcome	the impact upon an organization which is wholly or partially attributable to an **incident**.
Patient	a person who is a recipient of **healthcare**.
Patient characteristics	selected **attributes** of a **patient**.
Patient outcome	the impact upon a patient which is wholly or partially attributable to an **incident**.
Patient Safety	freedom, for a patient, from unnecessary harm or potential harm associated with **healthcare**.
Patient safety incident	an **event** or **circumstance** which could have resulted, or did result, in unnecessary **harm** to a **patient**.
Preventable	accepted by the community as avoidable in the particular set of circumstances.
Quality	the degree to which health services for individuals and populations increase the likelihood of desired health outcomes and are consistent with current professional knowledge.
Resilience	The degree to which a system continuously prevents, detects, mitigates or ameliorates **hazards** or **incidents**.
Risk	the probability that an **incident** will occur.
Root cause analysis	a systematic iterative process whereby the factors which contribute to an **incident** are identified by reconstructing the sequence of events and repeatedly asking why? until the underlying root causes have been elucidated.
Safety	freedom from **hazard**.
Semantic relationship	the way in which things (such as **classes** or **concepts**) are associated with each other on the basis of their meaning.
Side effect	a known effect, other than that primarily intended, related to the pharmacological properties of a medication.
Suffering	the experience of anything subjectively unpleasant.
System failure	a fault, breakdown or dysfunction within an organization's operational methods, processes or infrastructure.
System improvement	the result or outcome of the culture, processes, and structures that are directed toward the prevention of **system failure** and the improvement of **safety** and **quality**.
Violation	deliberate deviation from an operating procedure, standard or rules.

11 Appendix 2: Definitions from the AHRQ 'Common Formats'

Note: AHRQ uses asterisks (*) in the Common Formats glossary to indicate terms and definitions that are either identical (*) or similar (**) to the WHO ICPS preferred terms.

Accident	See: Patient safety incident.
Adverse event	See: Harm incident.
Adverse event outcome	See: Unexpected adverse outcome.
Adverse outcome	Undesired patient outcome of healthcare; clinical complication of healthcare (which may or may not be a patient safety incident).
Adverse reaction**	Unexpected adverse outcome resulting from a justified action where the correct process was followed for the context in which the event occurred.
Ameliorating action	See: Rehabilitation action.
Blocking action	An action taken to try to prevent an event such as a process failure or error from reaching a patient.
Bodily injury**	Physical harm or damage to a person's body.
Chain of events	A series of events in which one event leads to another with the possibility that the final event reaches the patient; collectively, an event episode.
Circumstance**	Condition surrounding and affecting a person, process, etc; as related to healthcare, the context within which processes are performed to deliver healthcare services, including culture of safety, management structure and incentives, staffing levels, qualification and training of staff, acquisition and maintenance of devices, condition of care environment, etc.
Close call	See: Near miss.
Complete event report	See: Complete patient safety concern report.
Complete patient safety concern report	A complete initial report of a patient safety concern, including, as applicable, data elements comprising the following Common Format paper forms: healthcare event report form, patient information form, one or more event specific forms, and, finally, final assessment form.
Complication	See: Adverse outcome.
Contributing factor*	A circumstance, action or influence which is thought to have played a part in the origin or development of an incident or to increase the risk of an incident.
Degree of harm*	The severity and duration of harm, and the treatment implications, that result from an incident.
Detection**	Identification of a patient safety concern; a broader concept than discovery as it also encompasses identification by automated and other systems.
Device	See: Medical device.
Discoverer	Person in a healthcare facility who identifies a patient safety concern. Such person should preferably either report the concern or inform another person in the healthcare organization of the concern (who then reports it).
Discovery	Identification of a patient safety concern during the course of performing duties in a healthcare facility; a narrower concept than detection.
Duration of harm	The period over which disease, disability, disfigurement, dysfunction, etc. may be evident; often denoted as none, transient, temporary (short-term) or permanent

	(life-long).
Error*	Failure to carry out a planned action as intended or application of an incorrect plan.
Error recovery	Recognition of and recovery from an error by the person making an error, thus averting an event or the creation of an unsafe condition.
Error recovery action	Action taken by a person to recover from an error that he or she made in an attempt to avert an event or the creation of an unsafe condition.
Event	See: Patient safety event.
Event chain	See: Chain of events.
Event episode	Single event or chain of events from point of origin of a process failure or error within a healthcare organization to its termination in a near miss or incident.
Event reporter	See: Reporter.
Facility	See: Healthcare facility
Factor	See: Contributing factor.
FAF	See: Final Assessment Form.
Fail-safe	Process designed to prevent the failure of a healthcare process or error made in the delivery of a healthcare service from propagating; usually an integral element of that process; may be a secondary system designed to insure the continued function or operation of a primary system.
Final Assessment Form (FAF)	Common Formats paper form that comprises data elements to assess and to encode a patient safety concern. The completion of this form completes the initial report of an unsafe condition, near miss, or incident. Such report may (or may not) be followed by an event episode analysis.
Harm	Physical or psychological injury (including increased anxiety), inconvenience (such as prolonged treatment), monetary loss, and/or social impact, etc. suffered by a person. See also narrower term Harm to health.
Harm from healthcare	See: Harm.
Harm incident*	An incident that resulted in harm to a patient.
Harm scale	See: Patient outcome harm scale.
Harm to health*	Impairment of structure or function of the body and/or any deleterious effect arising therefrom.
Healthcare-associated harm*	Harm arising from or associated with plans or actions taken during the provision of healthcare rather than an underlying disease or injury.
Healthcare Event Report Form (HERF)	Common Formats paper form comprising data elements designed to capture the essence of a patient safety concern.
Healthcare facility	Physical structure, such as a hospital, in which healthcare services are performed.
Healthcare location	Physical place within a healthcare facility; includes a location in which healthcare services are delivered (healthcare service delivery location), as well as such other areas as corridors, elevators, and those where supporting services are performed (e.g., laundry, meal preparation, and power generation).
Healthcare setting	See: Healthcare facility & Healthcare location.
HERF	See: Healthcare Event Report Form.
Incident	See: Patient safety incident.

Incident report	See: Healthcare event report form.
Injury	See: Bodily injury & Psychological injury.
Inpatient facility	Healthcare facility with beds for patients to stay overnight.
Location	See: Healthcare location.
Medical device	A medical device is any instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including any component, part, or accessory, intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, or intended to affect the structure or any function of the body, and which does not achieve its primary intended purposes through chemical action within or on the body and which is not dependent upon being metabolized for the achievement of its primary intended purposes.
Medical supply	See: Medical device.
Mitigating factor	See: Rescue action.
Multiple patient incident	An incident that involves two or more patients (e.g., patient given a medication intended for another - 2 patients affected by the same incident); includes incidents involving a population of patients, e.g., fire in a patient section of a hospital.
Near miss**	An event that did not, or could not, reach a patient. For example: discovery of a dispensing error by a nurse as part of the process of administering the medication to a patient (which if not discovered would have become an incident); discovery of a mislabeled specimen in a laboratory (which if not discovered might subsequently have resulted in an incident).
No-harm incident*	An incident which reached a patient but no discernable harm resulted.
No named patient event	See: Unsafe condition.
Origin of event	The initial process failure or error that eventually gave rise to a near miss or an incident; see also chain of events. An unsafe condition may have been a precursor to this initial failure or error.
Outcome	The result of a process acting on inputs in a given environment; as related to healthcare, often used to mean patient outcome.
Patient Information Form (PIF)	Common Formats paper form that comprises data elements to capture the characteristics of a patient involved in an incident.
Patient outcome	The result of receiving a healthcare service, especially as related to a patient's health status; usually refers to post-process results or measurements (the observed results of an intervention) whether or not one can confidently attribute those results to the preceding process (intervention).
Patient outcome harm scale	A systematic method to assess and to represent the extent of a patient's health loss (including anxiety, inconvenience, etc.) at a point in time, especially as related to residual harm following a patient safety incident and any rescue actions that might have been attempted consequently.
Patient safety accident	See: Patient safety incident.
Patient safety concern	Any circumstance involving patient safety; encompasses patient safety event (both incident and near miss) and unsafe condition.
Patient safety event**	Something that happens to or involves a patient; encompasses patient safety incident and near miss.
Patient safety incident**	A patient safety event that reached a patient, and either resulted in no harm (no harm incident) or harm (harm incident). The concept "reached a patient"

	encompasses any action by a healthcare practitioner or worker or healthcare circumstance that exposes a patient to harm. For example: if a nurse gives a patient an incorrect medication to take and the patient recognizes it as such and refuses to take it, an incident has occurred.
PIF	See: Patient Information Form.
Potential harm event	See: No harm incident & Near miss.
Preventable**	Accepted by the relevant community as avoidable in the particular set of circumstances.
Preventive measure	Process designed, or course of action taken, to keep something possible or probable from happening or existing; as related to patient safety, to prevent a patient safety event.
Process	A particular method of doing something, generally involving a number of steps or operations, that results in an outcome or produces an output.
Psychological injury	Harm or damage to a person's psyche, psychological functioning, or mental well-being.
Reaction	A patient's response to a stimulus or agent, such as to a medication.
Recovery	See: Error recovery.
Recovery action	See: Error recovery action.
Rehabilitation action**	Action taken, after any attempt to rescue a patient, to restore a patient's health status to what it was prior to an incident.
Reporter	Person in a heath care organization who reports a patient safety concern; may (or may not) be the person who discovered the concern.
Rescue action**	Action taken or started within the first 24 hours after the discovery of a patient safety incident that is intended to prevent, to minimize, or to reverse harm to the affected patient.
Residual harm	Estimated harm to a patient's health subsequent to any attempted rescue action taken or started within 24 hours after the discovery of an incident.
Resilience**	The ability to resist and/or to recover from process failures or errors; to continue to deliver (healthcare) services safely in the face of various faults and challenges encountered during the course of normal operations; includes the ability to anticipate and to adapt performance to current conditions.
Risk**	The product of 2 dimensions: 1) probability of an event's occurrence in a specified period, such as a year and 2) the harm that would typically result if the event were to occur; often limited to (1).
Safeguard	Aspect of a healthcare process that is designed to prevent harm from reaching a patient in the event of a failure or error (including the failure of a fail-safe).
Side effect	An effect (usually an adverse outcome) caused by something (such as a drug or procedure) that was not the intended or indicated effect. The occurrence of a known side effect, even if an adverse outcome, by itself, is not a patient safety incident. It should be considered a quality of care problem; if, for example, the prescribing physician failed to weigh properly the potential health benefits and risks of prescribing a medication with a potentially lethal, or otherwise adverse, side effect.
Substance administration	The process by which a substance (e.g., medication) is administered to a patient (whether or not the patient actually receives or takes it).

Surgical supply	See: Medical device.
Time of discovery	Date/time when a patient safety concern was discovered.
Time of occurrence of event	Date/time when a patient safety event occurred (point in time) or started (if it occurred over a period of time).
Unexpected adverse outcome	Adverse outcome that was not expected to be a result of the patient's treatment plan; harm suffered as a result of an incident.
Unintentionally retained item	Foreign object introduced into the body during a surgical operation or another invasive procedure, without removal prior to finishing the surgery or procedure, that the surgeon or other practitioner did not intend to leave in the body.
Unnecessary harm	Healthcare-associated harm that was not expected to result from a patient's treatment plan; harm resulting from an incident.
Unplanned intervention	An intervention that was not part of a patient's treatment plan prior to the event that necessitated the additional intervention.
Unsafe condition	Any circumstance that increases the probability of a patient safety event; includes a defective or deficient input to or environment of a care process that increases the risk of an unsafe act, care process failure or error, or patient safety event. An unsafe condition does not involve an identifiable patient. For example, an out-of-date medicine on a shelf represents an unsafe condition. It might be given to a patient, but the identity of such patient is unknown at the time of discovery. The attempt to administer the out-of-date medicine to a patient would either represent a near miss (if not administered) or an incident (if administered).

12 References

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

[2] Brown G, Carpenter R. Successful Application of Service-Oriented Architecture Across the Enterprise and Beyond. Intel® Technology Journal. 2004;8(4):345-60.

[3] Centers for Disease Control and Prevention. Public Health Conceptual Data Model - Premiere Edition. Atlanta 2000.

[4] World Health Organization. WHO Draft Guidelines for Adverse Event Reporting and Learning Systems. Geneva: WHO Document Production Services 2005.

[5] 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.

[6] Drafting Group of the Project to Develop the International Classification for Patient Safety. The Conceptual Framework for the International Classification for Patient Safety; Version 1.0 for Use in Field Testing 2007 - 2008. Geneva: World Health Organization; 2007 July.

[7] Datix Ltd. Software for Patient Safety and Risk Management. 2005 [cited 2008 April 10]; Available from: http://www.datix.co.uk

[8] National Patient Safety Agency. National Reporting and Learning System: an Overview. 2007 [cited 2008 April 10]; Available from: http://www.npsa.nhs.uk/patientsafety/reporting/general-overview-of-the-nrlsgeneral-overview-of-the-nrls/

[9] 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.

[10] Institute of Medicine. To Err is Human: Building a Safer Health System. Washington DC: National Academy Press 2000.

[11] 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

[12] 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.

[13] 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.

[14] Smith B. From Concepts to Clinical Reality: An Essay on the Benchmarking of Biomedical Terminologies. Journal of Biomedical Informatics. 2006;39(3):288-98.

[15] Jeong S, Kim H-G. Design of Semantically Interoperable Adverse Event Reporting Framework. *The Semantic Web – ASWC 2006*. Berlin / Heidelberg: Springer 2006:588-94.

[16] Arici S, Bertele P. ReMINE Deliverable D4.1: RAPS Taxonomy: Approach and Definition; 2008.

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

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

[19] 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.

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

[21] 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.

[22] Quality Interagency Coordination Task Force. Glossary of Terms. 2008 [cited 2008 April 14]; Available from: http://www.quic.gov/report/mederr8.htm

[23] 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.

[24] Institute of Medicine. Patient Safety: Achieving a New Standard of Care. Washington, DC: National Academy Press 2003.

[25] Edwards E. Man and machine: Systems for safety. *Proceedings of the British Airline Pilots Technical Symposium*. London: British Airline Pilots Association 1972:21-36.

[26] Molloy GJ, O'Boyle CA. The SHEL Model: A Useful Tool for Analyzing and Teaching the Contribution of Human Factors to Medical Error. Academic Medicine. 2005 February;80(2):152-5.

[27] World Health Organization. World Alliance for Patient Safety: Forward Programme 2005. Geneva: World Health Organization; 2004.

[28] 109th United States Congress FS. Patient Safety and Quality Improvement Act of 2005. Public Law 109-41 (S. 544). 2005.

[29] Agency for Healthcare Research and Quality - US Department of Health & Human Services. Paper Forms - AHRQ Common Formats for Patient Safety Organizations. 2008.

[30] Agency for Healthcare Research and Quality - US Department of Health & Human Services. Users Guide - AHRQ Common Formats for Patient Safety Organizations. 2008.

[31] Gangemi A, Guarino N, Masolo C, Oltramari A, Schneider L. Sweetening Ontologies with DOLCE *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management Ontologies and the Semantic Web* London, UK: Springer-Verlag 2002:166-81.

[32] IFOMIS. Basic Formal Ontology. 2006 December 2006 [cited 2007 January 25]; Available from: http://www.ifomis.uni-saarland.de/bfo/

[33] Ceusters W, Smith B. Strategies for Referent Tracking in Electronic Health Records. Journal of Biomedical Informatics. 2006 June;39(3):362-78.

[34] Ceusters W, Smith B. A Realism-Based Approach to the Evolution of Biomedical Ontologies. *Proceedings of AMIA 2006* 2006:121-5.

[35] The OBO Foundry. 2006 [cited Dec 2, 2008]; Available from: http://obofoundry.org/

[36] Smith B, Ashburner M, Ceusters W, Goldberg L, Mungall C, Shah N, et al. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. Nature Biotechnology. 2007;25:1251-5.

[37] Smith B. Ontology (Science). In: Eschenbach C, Grüninger M, eds. *Formal Ontology in Information Systems - Proceedings of the Fifth International Conference (FOIS 2008)*. Amsterdam: IOS Press 2008:21-35.

[38] Smith B, Ceusters W, Klagges B, Köhler J, Kumar A, Lomax J, et al. Relations in biomedical ontologies. Genome Biology. 2005;6(5):R46.

[39] Rosse C, Mejino JLV. A reference ontology for bioinformatics: The Foundational Model of Anatomy. Journal of Biomedical Informatics. 2003;36:478-500.

[40] Akkermans H, Gordijn J. Ontology Engineering, Scientific Method and the Research Agenda. In: Staab S, Svatek V, eds. *EKAW 2006, Lecture Notes in Artificial Intelligence 4248*. Heidelberg: Springer 2006:112-25.

[41] Ceusters W, Smith B, Goldberg L. A terminological and ontological analysis of the NCI Thesaurus. Methods of Information in Medicine. 2005;44:498-507.

[42] 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.

[43] 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.

[44] Johansson I. Bioinformatics and Biological Reality. Journal of Biomedical Informatics. 2006;39(3):274-87.

[45] Ceusters W, Smith B. Ontology and Medical Terminology: why Descriptions Logics are not enough. *Towards an Electronic Patient Record (TEPR 2003)*. San Antonio 2003.

[46] Spear AD. Ontology for the Twenty First Century: An Introduction with Recommendations. Saarbrucken, Germany: Institute for Formal Ontology and Medical Information Science 2006.

[47] 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.

[48] Ceusters W, Elkin P, Smith B. Negative Findings in Electronic Health Records and Biomedical Ontologies: A Realist Approach. International Journal of Medical Informatics. 2007 March;76:326-33.

[49] Smith B, Grenon P. The Cornucopia of Formal-Ontological Relations. Dialectica. 2004;58(3):279-96.

[50] Ceusters W, Manzoor S. How to track Absolutely Everyting? In: Obrst L, Ceusters W, Janssen T, eds. *Ontologies for Intelligence*. Amsterdam: IOS Press 2009 (in press).

[51] Ceusters W, Smith B. Referent Tracking for Digital Rights Management. International Journal of Metadata, Semantics and Ontologies. 2007;2(1):45-53.

[52] Ceusters W, Smith B. Referent Tracking for Corporate Memories. In: Rittgen P, ed. *Handbook of Ontologies for Business Interaction*. New York and London: Idea Group Publishing 2007:34-46.

[53] Manzoor S, Ceusters W, Rudnicki R. Implementation of a Referent Tracking System. International Journal of Healthcare Information Systems and Informatics. 2007;2(4):41-58.

[54] Ceusters W. Dealing with Mistakes in a Referent Tracking System. In: KS H, ed. *Proceedings of Ontology for the Intelligence Community 2007 (OIC-2007)*. Columbia MA 2007:5-8.

[55] Bittner T, Smith B. A Theory of Granular Partitions. In: Duckham M, Goodchild MF, Worboy MF, eds. *Foundations of Geographic Information Science*. London: Taylor & Francis Books 2003:117-51.

[56] Ceusters W. Towards A Realism-Based Metric for Quality Assurance in Ontology Matching. In: Bennett B, Fellbaum C, eds. *Formal Ontology in Information Systems*. Amsterdam: IOS Press 2006:321-32.

[57] Ceusters W. What is a diagnosis ? Workshop on Ontology of Diseases 2006 November 7 [cited 2008 December 4]; Available from: http://www.referent-tracking.com/RTU/papers/WhatIsaDiagnosis.ppt

[58] Scheuermann RH, Ceusters W, Smith B. Toward an Ontological Treatment of the Initiation, Manifestation, Recognition and Representation of Disease. 2009 (under review).

[59] 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

[60] European Committee for Standardization. EN 12388:2005. Health informatics - Time standards for healthcare specific problems. 2005.

[61] 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.

[62] Jansen L. Tendencies and other Realizables in Medical Information Sciences. The Monist. 2007 October 2007:(in press).

[63] McLaughlin S. The Diabetes Quality Improvement Project. Diabetes Spectrum. 2000;13(1):5-10.

[64] American Diabetes Association. Standards of Medical Care in Diabetes. Diabetes Care. 2005;28:S4-S36.

[65] National Committee for Quality Assurance. HEDIS 2008 Technical Specifications for Physician Measurement National Committee for Quality Assurance 2007.

[66] Ruttenberg A, Smith B, Ceusters W. Information-Artifact-Ontology. 2008 [cited 2008 December 10]; Available from: http://code.google.com/p/information-artifact-ontology/

[67] Ceusters W, Spackman KA, Smith B. Would SNOMED CT benefit from Realism-Based Ontology Evolution? In: Teich J, Suermondt J, Hripcsak C, editors. American Medical Informatics Association 2007 Annual Symposium Proceedings, Biomedical and Health Informatics: From Foundations to Applications to Policy; 2007 November 10-14; Chicago IL: American Medical Informatics Association; 2007. p. 105-9.

[68] Ceusters W. Applying Evolutionary Terminology Auditing to the Gene Ontology. Journal of Biomedical Informatics. 2009 (accepted)(Special Issue of the Journal of Biomedical Informatics on Auditing of Terminologies).