An Evolutionary Approach to the Representation of Adverse Events

Werner CEUSTERS [a,1], Maria CAPOLUPO [b], Barry SMITH [a], Georges DE MOOR [c]
[a] *New York State Center of Excellence in Bioinformatics & Life Sciences, Buffalo NY, USA*
[b] *Independent Consultant, Buffalo NY, USA*
[c] *Research in Advanced Medical Informatics and Telematics (RAMIT) vzw, University of Gent, Belgium*

Abstract. One way to detect, monitor and prevent adverse events with the help of Information Technology is by using ontologies capable of representing three levels of reality: what is the case, what is believed about reality, and what is represented. We report on how Basic Formal Ontology and Referent Tracking exhibit this capability and how they are used to develop an adverse event ontology and related data annotation scheme for the European ReMINE project.

Keywords. ontology, referent tracking, adverse events, patient safety

1. Introduction

The goal of the European Union funded ReMINE project (Grant Agreement 216134) is to develop a prediction, detection and monitoring platform for managing risks against patient safety (RAPS). This involved the development of an ontology by carrying out three types of activities. First was the description of the domain of adverse events as cognized by human beings. This formed the basis for a taxonomy organised in ways familiar to clinicians [1]. The taxonomy then served as input for the two other types of activities: the development of (1) a purpose-independent domain ontology for adverse events, and (2) a series of application ontologies derived from the former to support applications for guideline checking and protocol monitoring. In this paper, we report on our efforts to design the upper level parts of the domain ontology in such a way that it remains faithful to both the cognitive perspective reflected in the taxonomy and the principles for high quality ontology development as advanced under a realist agenda.

2. Objectives

The goal of the domain ontology is to support the management of a database populated through input from the ReMINE pilot sites that keeps track of incidents involving patients that can be qualified as 'adverse events'. The cognitive engineering position defended in [1] describes an adverse event as '*an incident (a perdurant) that occurred*

[1] Corresponding Author: Werner Ceusters, Ontology Research Group, New York State Center of Excellence in Bioinformatics & Life Sciences, University at Buffalo, 701 Ellicott street, Buffalo NY, 14203, USA; E-mail: ceusters@buffalo.edu.

to a patient during the past, that is documented in a database of adverse events and that is an expectation of some future happening that can be prevented' [2]. Our objective here is to show that by using the right ontological approach, a database built on a suitably modified version of this definition can not only be used to detect and monitor adverse events, but also to advance the state of the art towards the development of better mitigation and prevention strategies.

3. Material and Methods

We based our work on the ReMINE taxonomy as it was made available in August 2008 [1]. We further used the results of our analysis conducted on a number of distinct and mutually incompatible adverse event definitions found in the literature [3].

We followed the ontology development strategy proposed by the OBO-Foundry [4]. One Foundry requirement is 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 [5]. We used as feeder ontologies the *Basic Formal Ontology* (BFO) [6], the *Foundational Model of Anatomy* [7], and preliminary versions of two ontologies under development: the Disease and Diagnosis components of the *Ontology of Biomedical Investigations* [8] and the Information Artifact Ontology [9]. As further required by the Foundry, we used relations that are unambiguously defined according to the pattern set forth in the OBO Relation Ontology (RO) [10]. Finally, we provided for each representational unit a textual definition that can easily be translated into equivalent formal definitions. In light of the use of the domain ontology for the representation of adverse events that actually occurred and therefore are registered in a database, we also applied the principles of Referent Tracking (RT) [11]. Whereas BFO – and ontologies in general – focus on what is generic (*types*), RT focuses on what is specific (*instances*).

4. Results

Table 1 contains in alphabetical order those representational units (RUs) of the upper level part of the ReMINE ontology which are sufficient to represent the various configurations that may obtain between the occurrence of an incident, whether that incident involves some harm, whether it has been recognized as such by some cognitive agent, and whether it has been reported upon in the ReMINE adverse event database. For each RU we provide a term, a category, and a textual description.

The categories are formed on the basis of two criteria. The first concerns whether the specific entities – 'particulars' in BFO's terminology [12] – are a matter of (L1) first-order reality such as a disorder in a specific patient, (L2) beliefs or cognitive representations about portions of first-order reality on the side of a cognitive being, for instance a diagnosis formed by a clinician, or (L3) some durable representation of a first- or second-order entity (e.g., a statement in some record that the patient suffers from an allergy). The second criterion is whether the particulars are instances of a universal (indicated by 'U') or members of a defined class ('DC'), the latter being defined as a collection of particulars that are all instances of a universal (e.g., human being) and that share at least one characteristic (e.g., being located in Buffalo) that is also a characteristic of particulars that are not instances of that universal (e.g., houses in

Buffalo) whereby there are also instances of that universal that do not share that characteristic (e.g., human beings located in Sarajevo). For terms and phrases used in the descriptions, we used the following typographical conventions: (1) terms that express relationships are printed in **bold**, (2) terms borrowed from the ontologies mentioned in section 2 are in **SMALL CAPS**, (3) terms introduced in the ReMINE ontology are printed in *italics*.

5. Discussion

ReMINE's pragmatic definition for the term 'adverse event' provides useful information that can easily be understood by human beings, but that can not be formally defined in a way that allows unambiguous interpretation by software agents. This is for several reasons. First, there is nothing in reality that exhibits all the listed features simultaneously. Clearly, nothing that happened in the past (an occurrent in

[3]

Table 1. Upper level of the ReMINE domain ontology (version 1)

Term	Category	Description
act of care	L1/U	PROCESS (1) which **has agent** a *caregiver* and (2) as (passive) **participant** a *subject of care*, and (3) is **motivated by** an *underlying disease* and a *care intention*
adverse event	L1/DC	PROCESS **denoted-by** a *denotator* in a *RAPS repository*
adverse event denotator	L3/DC	*denotator* **denoting** an *adverse event*
caregiver	L1/DC	HUMAN BEING in which a *caregiver role* **inheres**
caregiver role	L1/U	ROLE **inhering** in a HUMAN BEING mandated to be the **agent of** *acts of care*
care intention	L2/DC	intention (DEPENDENT CONTINUANT) **inhering in** a *caregiver* that motivates him or her towards an *act of care*
denotator	L3/U	REPRESENTATIONAL UNIT **denoting** a UNIVERSAL or PARTICULAR
harm	L1/DC	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.
harm assessment	L1/U	PROCESS to determine whether another PROCESS is a *harm*
harm diagnosis	L2/DC	cognitive representation (DEPENDENT CONTINUANT) **resulting from** a *harm assessment* and involving a conclusion whether a certain PROCESS is a *harm*
history of a continuant	L1/U	PROCESSUAL ENTITY **formed by** PROCESSES in which a CONTINUANT **participates** or **has participated**
positive harm diagnosis	L2/DC	*harm diagnosis* involving a conclusion that a certain PROCESS is a *harm*
RAPS repository	L3/U	INFORMATION ARTIFACT **formed-by** *denotators* that *denote* PROCESSES that started **in relation to** an *act of care* and are the **subject of** a *positive harm diagnosis*
subject of care	L1/DC	HUMAN BEING **undergoing** *acts of care*
underlying disease	L1/DC	the DISEASE **inhering in** the *subject of care* which is **part of** what serves to motivate performance of the *act of care*

Table 2. Fragment of a RAPS repository

1. #3 **instance-of** *act of care*	9. #4 **inheres-in** #5 **since** t_3
2. #3 **has-participant** #2 **at** t_1	10. #5 **member-of** *caregiver* **since** t_6
3. #2 **member-of** *subject of care* **at** t_1	11. t_6 **earlier** t_1
4. #1 **instance-of** PROCESS	12. #7 **instance-of** *RAPS repository* **since** t_0
5. #1 **has-participant** #2 **at** t_2	13. #6 **instance-of** *adverse event denotator* **since** t_4
6. t_1 **earlier** t_2	14. #6 **part-of** #7 **since** t_4
7. #4 **member-of** *positive harm diagnosis* **since** t_3	15. #6 **is-about** #1
8. #4 **is-about** #1	16. #6 **has-author** #5

BFO's terminology, which corresponds to DOLCE's 'perdurant') can be an expectation (a cognitive continuant).

Second, the pragmatic definition does not explicitly specify what sorts of incidents are to be considered as adverse events. Under one reading, one could assume that an incident becomes an adverse event by the mere fact of being reported '*in a database of adverse events*'. But that – if applied incautiously – 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. Further questions raised by this definition are, accordingly: who has the authority to add reports, what criteria are used by that authority, how to deal with false positives and negatives?

For the envisioned database to be able to represent reality faithfully and to enable use of its data to advance the state of the art in adverse event management, the system should be set up in such a way that it can accommodate annotations for incidents separately from beliefs about whether these incidents are adverse events, and in such a way that adequate quality control measures are put into place. Thus suppose that an incident (#1) that happened at time t_2 to a patient (#2) after some intervention (#3 at t_1) is judged at t_3 to be an adverse event, thereby giving rise to a belief (#4) about #1 on the part of some person (#5, a caregiver as of time t_6). Then applying ReMINE's definition for adverse event requires the introduction (at t_4) of an entry (#6) to that effect in the database (#7, installed at t_0). Using the ontology of Table 1 and a syntax as in [10] expanded with the temporal representations standardised in CEN EN12388 [13], one can then make the assertions listed in Table 2.

This approach, which in contrast to related work reported in [3] provides an evolutionary view on reality, allows us to track in detail and with various kinds of subtleties how the relevant portions of reality and the stakeholders' beliefs therein evolve over time. Some subtleties are built into the ontology. As an example, criterion (3) for *act of care* excludes those processes whose agents are caregivers but which are not performed under the caregiver role (e.g., a doctor hurting a patient in a car accident on the parking lot of a care facility). Other subtleties come with the referent tracking approach exemplified in Table 2. So is clinician #5 in the described scenario careful (perhaps erroneously so) *not* to assert that '#1 **member-of** *harm*' (#17), although he could have done so. If his diagnosis, as expressed in assertion 7 in Table 2, is correct, then #17 would correspond with reality too. If, on the other hand, his diagnosis is incorrect, then the presence of #17 in the repository would be an error. However, such errors can be corrected at later stages without losing information about the original beliefs [14]. If RAPS repository #7 were faithful to reality, each member of *adverse event* would be a member of *harm*. Furthermore, if #7 were locally complete, each member of *harm* that occurred in the realm in which #7 is installed would be a member of *adverse event*. Many other assertions can be added expressing other beliefs about #1, or even beliefs about somebody else's beliefs. Distinct clinicians, depending on what

definition they apply, may indeed 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 their beliefs about what caused the incident, and about how to prevent future happenings of incidents of the same sort in the future. Moreover, they may change their beliefs over time.

6. Conclusion

We have used the principles underlying BFO and RT to develop an adverse event ontology and associated data annotation scheme for the ReMINE project. The three-layered structure of reality – what is the case, what is believed, and what is represented – as argued for in BFO, turns out to be essential in this domain. 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 an adverse event ontology to cover all three levels of reality, in addition to appropriate error management.

References

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

[2] Ceusters, W., Capolupo, M., Devlies, J. (2009) ReMINE Deliverable D4.2: RAPS Domain Ontology. RAMIT VZW, Gent.

[3] Ceusters, W., Capolupo, M., De Moor, G., Devlies, J. (2008) Introducing realist ontology for the representation of adverse events. In Eschenbach, C., Gruninger, M. (Eds.) Formal Ontology in Information Systems, IOS Press, Amsterdam, 237–250.

[4] Smith, B., Ashburner, M., Ceusters, W., Goldberg, L., Mungall, C., Shah, N. et al. (2007) The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology* 25:1251–1255.

[5] Smith, B. (2008) Ontology (Science). In Eschenbach, C., Grüninger, M., (Eds.) *Formal Ontology in Information Systems – Proceedings of FOIS 2008*, IOS Press, Amsterdam, 21–35.

[6] Grenon, P., Smith, B., Goldberg, L. (2004) Biodynamic ontology: Applying BFO in the biomedical domain. In Pisanelli, D.M. (Ed.) *Ontologies in Medicine*, IOS Press, Amsterdam, 20–38.

[7] Rosse, C., Mejino, J.L.V. (2003) A reference ontology for bioinformatics: The Foundational Model of Anatomy. *Journal of Biomedical Informatics* 36:478–500.

[8] Scheuermann, R.H., Ceusters, W., Smith, B. (2009) Toward an ontological treatment of disease and diagnosis. *Proceedings of the 2009 AMIA Summit on Translational Bioinformatics*, San Francisco, California, March 15-17, 2009;: 116-120.

[9] Ruttenberg, A., Smith, B., Ceusters, W. (2008) Information-Artifact-Ontology. http://code.google.com/p/information-artifact-ontology/

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

[11] Rudnicki, R., Ceusters, W., Manzoor, S., Smith, B. (2007) What particulars are referred to in EHR data? A case study in integrating referent tracking into an electronic health record application. In Teich, J.M., Suermondt, J., Hripcsak, G. (Eds.) *AMIA 2007 Annual Symposium Proceedings, Biomedical and Health Informatics: From Foundations to Applications to Policy*, Chicago, IL, 630–634.

[12] Smith, B., Kusnierczyk, W., Schober, D., Ceusters, W. (2006) Towards a reference terminology for ontology research and development in the biomedical domain. In *Proceedings of KR-MED 2006, Biomedical Ontology in Action*, Baltimore MD, USA, 57–65.

[13] European Committee for Standardization. (2005) *EN 12388:2005. Health Informatics – Time Standards for Healthcare Specific Problems.*

[14] Ceusters, W. (2007) Dealing with mistakes in a referent tracking system. In Hornsby, K.S. (Ed.) *Proceedings of Ontology for the Intelligence Community 2007 (OIC-2007)*, Columbia MA, 5–8.