Response to Reviewers[1]

of

AEO: a realism-based biomedical ontology for the representation of adverse events

Yongqun He, Zuoshuang Xiang, Sirarat Sarntivijai, Luca Toldo and Werner Ceusters

---------- REVIEW 1 ----------

Reading the abstract, the major advance in science in this paper is the introduction of the trivial fact that an ADE 'induced by' Medical Intervention, and the need to model ADE hypothesis.

→ The reviewer seems to be unfamiliar with the problem of what counts as 'adverse event' in mainstream approaches and the incomparability of data that results from this situation. It is exactly by making the explicit distinction between what we call 'adverse event' in contrast to 'adverse event hypothesis' that we turned a difficult problem into a 'trivial' one. That is, so we believe, indeed an advance in science. Einstein's '$E=mc^2$' would be trivial in the same sense as argued for by this reviewer, as it is now taught to high school kids.

At the end of the Introduction section, the authors state: "we present our current development of AEO, thereby distinguishing it from another recent effort to generate an Adverse Event Reporting Ontology (AERO).". But the comparison with AERO is restricted merely on a few lines in the Discussion section. A thorough comparison should be done in Sec 3.5. where ADR is compared with AEO.

→ These few lines describe the fundamental distinctions that there are. We fail to see what else (and non-trivial) there would be to say about it. Furthermore, AERO is under development, as we state, so comparison can only be done on what there is to compare with.

Section 3.1:
"Currently AEO has 484 representational units, annotated by means of 369 terms with specific AEO identifiers, and 115 terms imported from existing ontologies (Table 1)." --> It sounds strange to use the word 'annotation' like this.

→ Not if one is aware of the use-mention confusions that non-realist biomedical ontologies are suffering from. It would have been helpful if the reviewer would have given us his interpretation of the word 'annotation' such that we would have an idea how to phrase it better (for him). However, we welcomed this remark nevertheless because we were indeed sloppy in other parts of the paper by not using the terms 'representational unit' and 'term' appropriately. That has been corrected throughout.

AEO claims the realist perspective eases valid and sound logical reasoning. I would like to see this done in practice, i.e. show the added value of the inferred ontology vs the asserted. The

---

[1] Reviewers remarks are throughout this document printed in blue font. Responses are in black font and preceded by '→'

needed OWL 2 profile should be described as well and where they use DL constructors like existential vs universal quantifyers. (e.g. in Fig 3).

→ OWL and specifically DL are inadequate for representing that all continuants, in contrast to occurrents, exhibit relations in which time is one of the relata. Renderings of realism-based ontologies in OWL are simplifications which may suffice for specific purposes. But that is not something we want to 'hard-code' in our work.
Nevertheless, we do mention indeed that there is an OWL implementation, without however specifying how it relates to the vision described in the paper. We clarified this.

To my knowledge BFO is not axiomatised in its owl implementation and hence can not be used for DL reasoning. The authors should state why they did not use BioTop as a DL reasonable upper level.

→ First, this is not a paper about BFO. Second, biomedical ontology is not just about DL and covers more than pizzas, toppings and wine regions and as said before, focusing on DL is short-sighted. Third, there are so many things people have done differently but that does not force us to motivate our choices against each of these different ones.

Fig 1 could probably be merged into Fig 2, and self-standing nodes like Symptom, should be bound by edges to other nodes (i.e. AdverseEvent has_outcome Symptom).

→ No, this comment exemplifies exactly the sort of problems one runs into when looking to ontology from the pure computational science perspective. And it exemplifies also that this reviewer does not understand the real semantics of OWL-DL: a generic statement such as AdverseEvent has_outcome Symptom would mean that EVERY instance of AdverseEvent has SOME Symptom as outcome, and that is false given what the terms Adverse Event and Symptom denote under the framework proposed here.

Sec 3.2: Add the Manchester syntax for major log defined classes of Fig 3. Then also the differrence in "Can be" vs "has to be/have" will be clearer.

→ We are not in the OWL business. Further, we have no idea what the reviewer means here by 'can be' and 'has to be'. We startled by the thought that we would have used somewhere a relation such as 'may be' or 'can have'. We double checked and we didn't resort to this kind of nonsense at all.

I would introduce the key entities (sec 3.3) before the log def of AE (3.2).

→ That is a good suggestion which we thus followed

In Sec 3.3. you state "For some particular to qualify as an instance of adverse event as defined in AEO the following key entities must exist", implying the existential constructor usage for those relations embedding these classes.

→ The existential constructor in OWL is not the same existential constructor as typically used in logic in a general sense.

Fig 2 [now Fig. 1]: Add Symptoms to AE via a has_outcome link.

→ As explained above, this would be wrong.

Sec 3.5:
Put the promised comparison with AERO here.

→ We reorganized the paper

Sec 3.6: a PRACTICAL USE CASE should be the main section as the proof of the pudding is in the eating. But what this sec and Fig 3 presents is not a real use case,as the instances are artificial place holders. Here the pattern should be instantiated by a real data set in a real database. Show how you envision Data annotation / instantiation / class-individual assertion. I.e. who does the modeling/data annotation? At what time ? What data is likely to be present in Clinical Information Systems (CIS)? How is the formal semantics of your pattern exploited in an application?

→ We changed the title section to 'example' instead of use case

At the moment the realism based ontologies are still too theoretical.

→ This is nonsense. Do a thorough literature check and you will see that they are reported to be used in over 70 projects.

Show the practical value of your Foundry principles and realist perspective here in a concrete example. I like to see an A box and a Sparql query on that A box.

→ Figure 3 captures exactly in graphical format what is the case for the example given. SPARQL queries are only useful for people using SPARQL. We leave the translation to them.

Fig 3 is merely a graphical model in no standard exchange syntax. What are the little triangles? Blank nodes ? The object properties look distorted and have no constructors (a fact that makes me believe it was more OBO inspired?).

→ We agree that more clarification is needed and adapted the paper accordingly.

The displayes Instances are not taken from a real world case as one can see. Please use real world data and implement a link to your ontology. Why do you use RO:instance-of and RI:isa rather than the owl equivalents when developing in P4 anyway ?

→ Because we don't want to cripple our ontology by the shortcomings of OWL. At the other hand, we also don't want once again in this paper to put too much emphasis on these shortcomings which the last author of this paper as well as many others have published about

enough; for this time: live and let die. We welcome however that this reviewer would do so himself, publish it, and then we can comment on all the mistakes made (such as the AE - hasOutcome -symptom mistake). That will be very educational indeed.

What is the datatype for t1 and t2, and is it realistic to assume these can be drawn from a CIS ?

→ The reviewer confuses temporal regions with measurements of time and representations of the results of these measurements. But it is worth clarifying this indeed as other readers not familiar with the realist approach may be confused in the same way.

The section at the bottom of page 8 about the Brighton Collab should probably be moved to 3.5.

→ we re-arranged the paper

Sec 4: Rename:--> "Discussion and Conclusion" or better separate the discussion and conclusions.

→ we re-arranged the paper

This sec is too much another review of existing efforts, which should be put into the intro or Sec 3.5. Here is the place to discuss YOUR results.
I expect to see a discussion on the use case for AEO and Fig 3. Add nome quantifications to your results as well.

→ we re-arranged the structure of the paper to solve this issue.

Style:
Sec 3.2:
"Just temporal precedence is not enough"-->"Mere temporal precedence is not enough" OR "Temporal precedence alone is not enough"

→ this has been corrected

Sec 3.4:
"Rehan et al provides physicians guide to how to assess causal relations of adverse events induced by drug administration [11].
The causal requirement is different from that in concept of adverse event in existing adverse event reporting systems" -->
"Rehan et al provides physicians with a guide how to assess causal relations of adverse events induced by drug administration [11].
The causal requirement is different from that of the concept of adverse event in existing adverse event reporting systems"

→ This has been corrected

"The data stored in such an adverse event reporting system is typically used to generate hypotheses about whether there is causality involved between what is reported as adverse events and medical interventions." -->
"The data stored in such an adverse event reporting system is typically used to generate hypotheses on causal links between a reported adverse events and medical interventions."

→ we rephrased though not as proposed here because the term 'reported adverse event' is ambiguous.

Sec 3.5: Reword "Ceusters et al. provides an excellent survey and summary on different types of adverse event representation [9]" as it sounds too much like self-praising.

→ We did so.

---------- REVIEW 2 ----------
General comments:
The authors are to be commended for applying OBO foundry principles to the area of adverse events.
The ideas are presented with adequate clarity. The specification of a model for intervention-induced adverse events is a creative contribution.
The paper suffers somewhat from stilted or unnatural English, and could benefit from some careful editing in this regard.

→ We did so

The major deficit of the paper is lack of any attempt at evaluation of the model, and only oblique suggestions of the objective criteria for judging the value, applicability, or correctness of the particular design choices, outside the narrow context of OBO-based ontological analysis. As a workshop paper, this deficit is perhaps understandable, but a full paper would require a much better section on evaluation and comparison.

→ That is correct, but AEO is still work under development. The principles itself, specifically in light of the sort of mistakes these principles help to avoid, are worth reporting on. We clarified this in the paper.

The paper makes some rather bold assertions about what "adverse event" really ought to mean, …

→ That is not our intention. Thus we made it more clear that it reflects OUR use of the term within AEO.

… when many clinical groups have already attempted to establish that for their purposes a firm causal relation is not required for reporting adverse events. It would seem prudent and less overtly dismissive of prior work to use the term "intervention-induced adverse event" for the meaning the authors intend in this paper. It would seem highly unlikely that practitioners will

abandon their more liberal use of the term "adverse event" based on the authors' nice model; rather, the authors could advance the cause of clarity of meaning by adopting a term that clearly communicates their meaning.

→ that would be acceptable if the field had come to one accepted definition for 'adverse event'. But as shown in one of the papers referenced in our paper, that is not the case.

A few minor comments:
Abstract: the number of terms is listed as 384, but this disagrees with the paper, which says 484.
dosis -> dose
un-wanted -> unwanted

→ all corrected

---------- REVIEW 3 ----------
This paper gives a interesting illustration of the building and use of an adverse event ontology. I have a couple of remarks that illustrates more a different view from my point than a fundamental criticism on the quality of the paper. Still, it may enhance the quality of the paper and reduce unnecessary reactions if these remarks would be considered.

- the problem of importing sub-ontologies or parts of sub-ontologies. To me this looks like a non scalable approach. Once imported it becomes part of your ontology and then you are responsible for it. If the other ontology evolves you loose connection. So I would not call this "... supports efficient ontology reuse ..." (page 5 by the way only the even pages are numbered). I realize however that this is a a topic in its own right.

→ the reviewer is right for mainstream ontology development, but not in our case because we make exclusively use of ontologies following the principles of Ontological Realism and these principles have been designed to avoid such problems. We made this more clear in the paper.

- page 5 a clinically abnormal process (e.g. a fever process) = a dangerous example. Is fever abnormal? What are the normal boundaries ? Does it depend on age, gender, ... ? Part of the adverse event definition in real life will be a clinician saying " this is not normal". So, maybe adverse event is not a subclass of pathological bodily proces but rather of bodily process deemed pathologically by someone.

→ the example is not dangerous, but clearly not well explained enough in the submitted version. This has now been corrected along the following lines: whether an instance of bodily process is pathological or not, is not a matter of opinion, but of reality. But observers observing that instance might come to the wrong conclusion. That is exactly the reason why adverse events need to be distinguished from adverse events hypotheses.

My point is: maybe there exists something like pathological bodily proces, but what's the point if, in operational systems you need more concepts to come to that. I assume that the ultimate goal of this effort is to "make a system", right ?

→ The goal is in the first place to have a faithful description of the entities and relationships that obtain between them in the domain covered. Operational systems should then select from these descriptions what is relevant for their purposes. We highlighted that too.

At least for operational systems I favour an approach where you don't try to make any possible combination of concepts and make then new concepts of it, giving it an identifier and a name. This tends to create an overwhelming amount of concepts and unuseable and complex ontologies.

→ Realism-based ontologies are as complex as reality (as far as we know wrt the state of the art). It is then always possible to simplify for specific purposes.

A solution to that is the use of blank nodes, where you limit the concepts but you group them together using blank nodes. (much like postcoordination in coding systems).

→ blank nodes are an RDF oddity we don't want our original representations to suffer from. How one wants to group data - this is what RDF blank nodes are for - has nothing to do with what entities are and how they behave in reality.

The same thing holds for adverse event versus adverse event hypothesis , where to me, adverse event hypothesis is overspecified. (but this is not part of AEO, so I cannot blame you for this :-)

→ we don't see what is the problem with 'adverse event hypothesis' at least if the 'adverse event' part denotes what AEO denotes by 'adverse event'.

- what about " fever as a clinical abnormal proces (page 5) versus fever as a symptom (page 6)? this may be confusing unless explaining that a symptom can be an abnormal proces or vice versa or there is some relation between them.

→ agreed that there is confusion possible for readers not familial with OGMS. We removed the sources for confusion.

These were just thoughts and they didn't affect my score, because it merely reflects different approaches.