

Response to reviewers for the paper  
***Representing SNOMED CT Concept Evolutions using Process Profiles***  
*W. Ceusters and J. Bona*

This paper was submitted originally to the FOIS 2016 main conference for which it was rejected and then in its original form – eliminating only some figures to accommodate distinct page restrictions – to the International Workshop on Ontology Modularity, Contextuality, and Evolution (WOMoCoE 2016), part of FOIS 2016, since it addresses the topics suggested there. For the final paper published in the Workshop's proceedings we addressed the reviewers' concerns of both the main conference and the workshop where reasonable.

**Reviews from WOMoCoE 2016**

Dear Werner,

We are happy to inform you that your paper '1: Representing SNOMED CT Concept Evolutions using Process Profiles' has been accepted for WOMoCoE 2016.

Reviews for your submission are included below: please take carefully into account the reviewers comments in preparing the final version of your paper.

Please be reminded that the deadline for camera ready submission is: 15 June 2016.

We will shortly provide detailed information for the camera ready submission and instructions for the presentation of your work at the workshop.

Kind regards,

WOMoCoE 2016 PC chairs

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

OVERALL EVALUATION: 2 (accept)

REVIEWER'S CONFIDENCE: 4 (high)

Relevance: 5 (excellent)

Potential: 3 (fair)

Quality: 4 (good)

In this paper, the authors describe the features of Snomed's Release Format 2 (RF2) to describe the changes made in every released version, and propose a uniform representation to evaluate these changes, and the whole ontology, through process profiles.

I find this to be a pretty good example of a workshop paper. The authors identify a specific problem that they want to tackle, explain it in detail, propose their solution, and analyse its properties. Although the specific contribution is relatively minor, and only applicable to this specific case (RF2 is only used for Snomed releases), it provides an interesting use case where tracking the evolution of knowledge (being them axioms or consequences) is of utmost relevance.

**→ Thanks for this comment. Our work is currently indeed of limited applicability not because RF2 is only used by SNOMED CT, but because almost all other ontologies lack a formal mechanism for communicating change at all. I highlighted this in the paper.**

I believe that the paper should be accepted.

----- REVIEW 2 -----

OVERALL EVALUATION: 0 (borderline paper)

REVIEWER'S CONFIDENCE: 4 (high)

Relevance: 4 (good)

Potential: 2 (poor)

Quality: 2 (poor)

The paper addresses the representation of changes of the SNOMED CT ontology over time. The aim is to assess "how faithful" the changes are in respect of the biomedical reality". The methodology to achieve this aim and how to measure progress ("use reality as a benchmark") towards this aim stay unclear.

**→ That was NOT the aim of the paper, rather the aim of our past efforts as clearly stated in the introduction. Both the methodology and how to measure progress are explained at length – also as stated in the introduction – in the following references:**

**[4]W. Ceusters, and B. Smith, "A Realism-Based Approach to the Evolution of Biomedical Ontologies," *Biomedical and Health Informatics: Proceedings of the 2006 AMIA Annual Symposium*, pp. 121-125, Washington DC: American Medical Informatics Association, 2006.**

**[5]W. Ceusters, "Applying Evolutionary Terminology Auditing to SNOMED CT," *AMIA Annu Symp Proc*, vol. 2010, pp. 96-100, 2010.**

**[6]W. Ceusters, K. A. Spackman, and B. Smith, "Would SNOMED CT benefit from Realism-Based Ontology Evolution?." pp. 105-109.**

The paper is very non-technical (in the sense that it presents an abstract, scientific view on the matter), but resorts to example-base narrative. More than roughly 3/4 of the paper are a description of the information that is captured in the RF2 format. The aim, contribution and improvement by the suggested methods stay unclear to the reader. Furthermore, it is entirely unclear how the observed effects and the envisioned aim can be achieved for other ontologies.

**→ Again this reviewer is in error: 2.5 pages describe the relevant aspects of SNOMED CT the way it is – a description which is needed to make the paper understandable on its own – while the rest deals either with the basics of the theories we use, and the transformation of the information conveyed by means of RF2 into process profiles. No actions taken in relation to this comment.**

Thus this paper is too narrow in coverage and too immature to be presented at a scientific workshop.

**→ I believe this reviewer's comments to be too narrow and immature, specifically in light of his remark that this paper is non-technical and scientific. It is new to me that being scientific is frowned upon at a scientific conference.**

----- REVIEW 3 -----

OVERALL EVALUATION: 2 (accept)

REVIEWER'S CONFIDENCE: 4 (high)

Relevance: 5 (excellent)

Potential: 5 (excellent)

Quality: 4 (good)

The paper starts with an overview of the modelling and philosophical foundations behind SNOMED CT RF2, the latest release of one of the largest and prominent healthcare terminology currently available. The paper continues with a description of the approach implemented to represent the changes introduced from previous versions of SNOMED CT to version RF2.

SNOMED CT represents one of the most outstanding contributions in the field of ontology engineering. For nearly 20 years (more if we consider SNOP as the starting point), SNOMED CT has been refined and extended to reach a level of maturity difficult to match by related initiatives. As a result, this ontology has been used in many biomedical applications.

**→ we agree**

In the first sections of the paper, the authors introduce some of the modelling and philosophical foundations considering during the design of SNOMED CT RF2. According to the authors of the paper, "an ontology is a faithful representation of the part of reality that it covers".

**→ The introduction presents the foundations of the theories according to which the authors – and the community of 'realist ontologists' – measure the quality of ontologies. SNOMED CT has not been designed at all with these foundations in mind. It is the case however that it satisfies certain principles thereof and that there is a clear tendency towards moving further in that direction.**

The authors illustrates this notion of ontology with the holographic simulation of an aquarium. The information included in the holographic simulation is so precise that it would even be possible to predict a future state of the aquarium. Although it is not mentioned explicitly, it gives the impression that the authors assume that there is only one reality and that this reality can be faithfully modelled.

**→ That is exactly what we entertain as view.**

I have the opinion that several interpretations of the reality are possible. Our understanding of a phenomenon can also change in the course of time. For instance in Physics, the notions of gravity and matter have quickly evolved in the past decades. In addition, we have to consider the point of view of the ontology engineer and the possible applications of a particular ontology. For instance, the properties of a bridge that are of interest of an engineer might be different from a worker that participates in the construction of the bridge. This is why building an ontology is so difficult, in particular, if you try to satisfy the requirements of a broad audience.

**→ The existence of only one reality does not exclude the possibility of there being distinct interpretations about. As knowledge accumulates, our ontologies change indeed, but not the**

*reality up to that point in time. 'Faithful' does also not mean 'complete' although that is what should be aimed for by reference ontologies, in contrast to application ontologies. These aspects are covered in, f.i., [2] B. Smith, and W. Ceusters, "Ontological realism: A methodology for coordinated evolution of scientific ontologies," Applied Ontology, vol. 5, no. 3-4, pp. 139-188, 2010.*

SNOMED CT RF2 was created following this realism-based perspective and supported by the BFO and IAO ontologies.

**→ As far as we know, that is not the case, but once again, we do note – and very much welcome – a tendency to move towards such perspective.**

The authors of the paper also include some details on how the components (concepts, relationships and descriptions) have been defined. From a philosophical point of view, I can agree more or less on this realism-based perspective but from a research point of view, I must acknowledge the depth of this contribution supported by several previous papers listed in the bibliography. SNOMED CT is not only based on common sense and the experience gathered during nearly 20 years of modelling and collecting feedback from many practitioners but also in some serious thought about the nature of ontologies and what they should represent.

In Section 2 the authors discuss the notion of concept in SNOMED CT. Each concept has a unique identifier and several associated descriptions. In particular, it is discussed the FSN that might include a "semantic tag" for disambiguation purposes. I understand the motivation behind FSNs and their "semantic tags" but given the importance of FSNs in applications of SNOMED CT, I would prefer a different approach where "semantic tags" are independent descriptions associated to the FSN of each concept. To identify a concept, we could also compose the FSN and the related semantic tags. Semantic tags seems to be influenced by modelling decisions in SNOMED CT. A modelling change that affect one of these semantic tags implies that many FSNs might be affected too.

**→ Indeed. However, we are not sure whether this reviewer is aware of our relation to SNOMED CT. For clarity, the authors of this paper ARE NOT AUTHORS of SNOMED CT. We rather use SNOMED CT as an object of study which leads to recommendations which the SNOMED CT authors then chose to implement or not. We absolutely agree that either a more formal treatment of semantic tags would be in place or that the idea should be abandoned at all because of the idiosyncratic relationship with high level concepts.**

In Section 2, it is also discussed the notions of "inside" and "outside" of the ontology that it is not clear to me. I will appreciate a bit of more detail what you really means with this.

**→ we elaborated on the example we provided already a bit more.**

As a general comment for the paper, I missed that examples are not presented using also DL notation. Given that the syntax and semantics in DL are well-defined, it would facilitate the understanding of what you really means in the examples.

**→ we used the syntax – though a bit simplified – offered by SNOMED CT itself.**

I also got the impression after reading the paper that the DL classifier is used after changes in SNOMED CT has been implemented and not during the implementation of the changes. I would appreciate if the authors can clarify when the DL classifier is used and in case that it is only used after changes are implemented, I would like to know why.

**→ we are not the authors of SNOMED CT. The procedure is explained in the SNOMED CT documentation. Page restrictions prevent us from elaborating on this in the paper.**

To preserve backwards compatibility with applications implemented using previous versions of SNOMED CT, components are persistent (never disappear in future releases and their identifiers are never reused) and they can be set as “active” and “inactive”. This implies that components continue to be distributed even when they are no longer active. I am fine with the idea that identifiers are never reused but keeping inactive components inside of the ontology does not seem a good idea. After several releases, it can be really messy to have active and inactive components. Most DL classifier will not distinguish between both types of components and they can produce undesired inferences. For practitioners would also be difficult to inspect the ontology using some of the available ontology editors and they might not be able to prevent unexpected effects for the presence of inactive components. I was just wondering if it would not be better to distribute SNOMED CT with only active components and provide a diff-ontology that represents the changes implemented from previous releases. Together with the diff-ontology, it can be included some mapping rules that explicitly indicate certain correspondences between the latest release of SNOMED CT and the diff-ontology.

**→ That diffs are not sufficient is part of the point we tried to make. Take a well-known example from ICD. ‘Nicotine dependence’ was originally not in ICD, not because it did not exist, but because nobody knew it existed. Then it appeared in the classification. A diff would only tell you that it is now there. But the point is: why? Did nicotine dependence not exist prior to its inclusion in ICD? Or did it exist, but scientists of these days didn’t know it existed? A diff does not tell you that. Then, later ‘nicotine dependence’ disappeared. A diff would tell you that. It would not tell you that it disappeared because of lobbying by the tobacco industry. Afterwards, until now, it is again there. See the picture?**

Section 5 requires more detail a clarity about what is presented. For instance, “each PPR consist of 29 characters, each one representing the status of some quality-like features that can be ascribed to the concept”. Later in the same section the authors wrote: ““A” in this case stands for active, while “\_” means that there is at the respective time no instance of the quality-like feature inhering in the concept. C4, in contrast, exhibits a different PPR for this feature, one that is the result of a start in the 9th version.”. I can deduce that the 29 characters in a PPR represents 29 versions of SNOMED CT, right? What is the meaning of “D” and “L”? Might be “D” corresponds to “dying” and “L” to “limited”?

**→ we added the explanation to the legend of Table 7.**

After reading Sections 4 and 5, I was wondering what is really inside (in terms of files and documentation) of a distribution of SNOMED CT.

**→ that is described in the 757 page SNOMED CT technical implementation manual provided as reference [10]. Hard to provide a perfect summary within the page limits set forward for this workshop.**

Section 6 seems to be very poor in terms of discussing related work. There is a lot of relevant work on change management and ontology evolution that it is not discussed in this section. I would definitely like to know more about which other work has been considered and why the authors have discarded several relevant approaches and proposed their own solution to manage changes on future releases of SNOMED CT.

**→ we referenced 9 relevant papers describing related work. More can be found in references 4 and 5. We challenge this reviewer to find related work that addresses the issues we described, and exemplified further here in this response by means of the nicotine dependence terms in ICD.**

I think the reason of lack of details in Sections 5 and 6 is that the paper is very ambitious from the content point of view. It is very difficult to discuss several aspects of the modelling and change management of SNOMED CT in a paper of only 12 pages (including references).

→ ***You bet.***

I think it would be better to focus on Sections 4-6 and discuss philosophical and modelling foundations in a separated paper. To present Sections 4-6, it would be enough to summarize some of the notions introduced in previous sections.

→ ***The sad news is that the foundations are essential, but rarely well-enough understood, in the first place by computer scientists and software engineers: they typically exhibit an 'inside' looking view, thereby making no relationship with the outside.***

It is a pity that I cannot attend to the workshop this year because I would love to discuss with the authors of the paper some of the issues I mentioned in the review.

→ ***feel free to contact us***

----- REVIEW 4 -----

OVERALL EVALUATION: 0 (borderline paper)

REVIEWER'S CONFIDENCE: 4 (high)

Relevance: 4 (good)

Potential: 4 (good)

Quality: 3 (fair)

This paper describes an attempt to formalise the evolution of the SNOMED CT terminology using the BFO-based ideas. The authors discuss the proposals for the Information Artifact Ontology that provide a realism-based perspective of the 'representation' relationship and apply these ideas to the changes of the SNOMED CT.

The authors stated that an ontology is a faithful representation of the part of reality that it covers.

→ ***Our view is that it should be. Unfortunately, very few really are.***

Thus changes in the ontology should correspond to changes in the reality. The authors first describe a SNOMED-CT terminology as a concept-based ontology, then represent main SNOMED CT concepts using the Information Artifact Ontology. In the next section the authors present a SNOMED CT changes format and provide their own definition of SNOMED CT changes based on history.

After reading a paper I am not convinced that the proposed format is suitable replacement/addition to the existing SNOMED CT changes format.

→ ***Why not?***

The authors starting from the introduction emphasise the need of a correspondence between an ontology and reality, calling it faithful representation. However, in the SNOMED CT changes formalisation, they specifically outline that only internal elements (corresponding to ontology rather than the real world) would be studied.

**→ No, we rather state explicitly that SNOMED CT provides only information over internal changes, and deplore that only vague descriptions of certain reasons are provided. We in contrast try to indicate how such internal changes relate to external changes.**

The main results are summarised in the Table 7, however the presentation of these results could be improved. "Each PPR consist of 29 characters..." -- does this mean that each character correspond a single SNOMED CT release? this is not stated in the paper.

**→ The assumption is correct and we highlighted this now explicitly.**

The legend for characters in the rows in a grey background ('D', 'L', 'P') is absent.

**→ We added this to the legend.**

Most importantly, the paper does not answer how this additional representation can be used. How can different concepts/change sets be compared using new history information? What is an additional value over the RF2 change format? This is not clear from the paper at all.

**→ The advantage is in the standardization and explicitation, obviously.**

This paper also lacks conclusion/Discussion section which will summarise the results of the paper.

**→ Discussion and conclusion were in the paper, though not explicitly labeled as such. We did now.**

Minor remarks and typos:

- Page 1, Abstract: "Its new release format..." -- It's

**→ Not quite. See <http://grammarist.com/spelling/its-its/>**

- Page 3: "...ontologies in general –to be..." -- spacing around '-'

**→ corrected**

- Page 4, T1 definition: '-' missing between attribute and object

**→ corrected**

- Page 5, Table 1, D6: "specifically depends on" missing underscores, that are used later (in E2)

**→ corrected**

- page 6, last paragraph: "...is also used as argument..." -- as an argument

**→ corrected**

- page 7, Table 3, Legend: "Active: (1) = active, A(0) = inactive." --remove A. Also several times later

**→ corrected**

- page 9, last paragraph: "...at a level on a a par with..." -- on a par with

**→ corrected**

- page 10, Table 7. I suggest to sort the attributes in all Ci for the ease of reading. E.g., for every concept start with FSN, followed by Is-a (where applicable), followed by Same-as, and then Was-a.

**→ changed as requested**

## Reviews from FOIS 2016

Dear Werner Ceusters,

Thank you for your submission to FOIS 2016. We regret to inform you that your paper: Uniform Representation of SNOMED CT Components using Process Profiles: a Matter of Life and Death has not been accepted for presentation at the conference.

Copies of the reviewers' reports on your paper are appended below, we hope these will explain the decision. We understand that every paper is the result of careful thought and hard work of its authors. We hope that the reviewers' comments will be useful in your research.

While your paper was not accepted for FOIS 2016, you may still consider to submit a revised version, taking into account the reviewers' comments, to one of the workshops:

In case one of the authors of your paper is a master student, PhD student, or postdoc, they have the opportunity to submit a short version of this (or other) work to the Early Career Symposium.

Note that the early deadline to register for FOIS is June 5 2016.

Sincerely,

Roberta Ferrario and Werner Kuhn

FOIS 2016 Programme Chairs

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

OVERALL EVALUATION: -1 (weak reject)

REVIEWER'S CONFIDENCE: 4 (high)

Scientific or technical quality: 2 (fair (bottom 1/3))

Novelty or innovation: 2 (similar to other work but still somewhat innovative)

Presentation: 3 (good)

References: 3 (nothing missing but irrelevant references present)

Recommendation for FOIS Best Paper Award: no

This paper presents some work into an investigation looking at historical changes in SNOMED-CT. The ideas and thoughts presented in the paper are nascent, but interesting. The long term research objective is particularly compelling.

The paper reads well, although at times appears to simply serve as a criticism of SNOMED and its developers. Furthermore, a lot of the presentation, although slick and polished, essentially feels like lab book notes and documentation on RF2. I think this latter point is simply due to the preliminary nature of the work.

***→ It is rather because of an attempt to summarize 757 pages of SNOMED CT's technical implementation manual so that readers of our paper don't have to dive into that document. In a journal paper we would benefit from the absence of page restrictions.***



My main issue with the work is that it comes over as being very preliminary. There is a lot of background information, which is well presented and informative, but there are no significant experiments, research results or conclusions. At this stage, the work seems more appropriate for a workshop rather than a conference like FOIS.

→ *And so we did.*

----- REVIEW 2 -----

OVERALL EVALUATION: -2 (reject)

REVIEWER'S CONFIDENCE: 5 (expert)

Scientific or technical quality: 4 (very good (upper 1/3))

Novelty or innovation: 3 (innovative)

Presentation: 2 (needs minor improvements - give details in the Review section below)

References: 4 (excellent references)

Recommendation for FOIS Best Paper Award: no

This paper proposes a new approach to address versioning in SNOMED CT. More specifically, the authors propose to use "process profiles", which corresponds to serializing the active status of any SNOMED CT component as a matrix of the type element x version.

Pros: 1) generally interesting and in scope for FOIS, 2) innovative solution, 3) versioning of ontologies is still an unsolved problem, 4) proposes a different serialization for the activity status of SNOMED CT components

Cons: 1) the benefit of using this approach is not demonstrated, 2) conclusions are speculative and not entirely substantiated, 3) unclear state of implementation, resources required, and scalability to the entirety of SNOMED CT, 4) unfocused paper with multiple arguments (SNOMED CT as an information content entity seems to be orthogonal to versioning -- or the link between the two is not well articulated), 5) limited significance: The paper essentially proposes only a different serialization for the activity status of SNOMED CT components.

In its current state, this preliminary paper would be more appropriate as a poster.

Additional comments

- subtitle is unnecessarily dramatic, especially given the lack of substantiation

- paper is not easy to follow, due to many acronyms, use of identifiers, etc.

- effectual should be defined early on

- the first paragraph of p. 7 is just wrong and is contradicted later on. The activity status for a component can be determined unambiguously for a given date.

→ ***Here the reviewer is wrong. We found that it CAN NOT be determined as such without computing for each version a complete transitive closure. See also [12] W. Ceusters, "SNOMED CT revisions and coded data repositories: when to upgrade?," AMIA Annu Symp Proc, vol. 2011, pp. 197-206, 2011. It is thus for this reason that our proposed work is more than just another serialization.***

Typos

- Each PPR consist\*s\*
  - ref. 6 is incomplete
- **corrected.**

----- REVIEW 3 -----

OVERALL EVALUATION: -2 (reject)  
REVIEWER'S CONFIDENCE: 4 (high)  
Scientific or technical quality: 2 (fair (bottom 1/3))  
Novelty or innovation: 2 (similar to other work but still somewhat innovative)  
Presentation: 1 (needs major improvements - give details in the Review section below)  
References: 2 (important references missing - give details in Review section below)  
Recommendation for FOIS Best Paper Award: no

The paper takes an interesting direction of work and proposes to infer changes in reality from changes in ontologies.

→ ***Not quite. We demonstrated in the past – as documented by the many self-citations, ‘self’ because nobody else seems to be interested in doing research in this topic – that where it should be the case that such inferences can be made, it cannot be done currently due to the simplicity and inadequacy of representation languages for that purpose, and the lack of information provided by ontology authors. SNOMED CT is a notable exception that goes in the right direction, though not enough (yet).***

However, it does so in a very preliminary and essayistic form. It embarks on wide ranging considerations, including the old debate of what concepts are, without making much progress on them, so that it ends up lacking focus and clarity of argument.

→ ***This reviewer does not see the progress as he clearly underestimates the problems with the concept-based approach.***

No clear research question is stated or implicitly addressed. The questions that are formulated are implicit ("explore") or vague and they seem somewhat different in each iteration they are stated.

→ ***The perception of the reviewer is, again, based on an insufficient understanding of the science behind our work. Our focus is clear, as stated explicitly: “Here we look instead at mechanisms that an ontology can offer to let us see changes in reality in a reliable way by examining the changes in the ontology”***

The paper ends without conclusions, which makes it unpublishable, as it is not clear what conclusions the authors could state if they were asked to add them.

→ ***This is wrong: the conclusions were in the last paragraph of section 6. We added now an explicit header.***

The references contain way too many self citations.

→ ***Reading the papers might have improved this reviewer’s understanding.***

Not directly linked to my negative assessment of the paper (which has to do with its immature state and incomplete form): The somewhat naive realism that seems to underly the work (namely, that there is an objective reality out there to be observed and recognized ...

**→ We are sure it is more naïve to believe there is no such objective reality, or that reality is not 'objective' in the sense described.**

... in the same terms that an ontology uses, no matter whether that ontology is shared by the observer or not) would seem to be hurting the undertaking. An ontology "knows" nothing about reality that it has not been told. And those who tell have a certain view of reality that may be more or less compatible with what the ontology contains already. Consequently, no amount of ontology analysis can reveal anything about reality and its changes per se, ...

**→ it cannot do so currently for reasons explained in our first response to this reviewer, and precisely therefore, the way ontologies are distributed should be changed such that it would become possible.**

... only about how such changes have been captured by humans.

Furthermore, it is impossible to distinguish whether ontology modifications result from observed changes in reality or changes in understanding of reality.

**→ Same point. Nevertheless, we have provided a mechanism to do so in earlier work and demonstrated that it can be used to predict how the ontology changes in the future. See [4] W. Ceusters, and B. Smith, "A Realism-Based Approach to the Evolution of Biomedical Ontologies," *Biomedical and Health Informatics: Proceedings of the 2006 AMIA Annual Symposium*, pp. 121-125, Washington DC: American Medical Informatics Association, 2006. [5] W. Ceusters, "Applying Evolutionary Terminology Auditing to SNOMED CT," *AMIA Annu Symp Proc*, vol. 2010, pp. 96-100, 2010. [8] W. Ceusters, "Applying Evolutionary Terminology Auditing to the Gene Ontology," *Journal of Biomedical Informatics; Special Issue of the Journal of Biomedical Informatics on Auditing of Terminologies*, vol. 42, no. 3, pp. 518-529, 2009.**

These fundamental difficulties need to be addressed with a clearer notion of 'concepts' and their relation to reality than the paper does. The paper does embark on some of this clarification, but does not yet reach closure or insights worth publishing.

**→ We beg to differ.**