

Scrutinizing the relationships between SNOMED CT concepts and semantic tags

Jonathan Bona^{1,*} and Werner Ceusters²

¹ Department of Biomedical Informatics, University of Arkansas for Medical Sciences, Little Rock, USA

² Department of Biomedical Informatics, University at Buffalo, Buffalo, NY

ABSTRACT

The *fully specified name* of a concept in SNOMED CT is formed by a *term* to which in the typical case is added a *semantic tag* (ST). An ST is meant to disambiguate homonymous terms and indicate where that concept fits into SNOMED's massive concept hierarchy. We have developed a method to determine whether or not a concept's tag correctly identifies its place in the hierarchy, and applied this method to an analysis of all active concepts in every SNOMED CT release from January 2003 to January 2017. Our results show that there are concepts in every release whose tags do not match their placement in the hierarchy. These tag/hierarchy mismatches appear to be errors. The number of such errors is increasing in recent versions.

1 INTRODUCTION

SNOMED CT is a large reference terminology for the clinical domain made up of 300,000+ active concepts with machine-readable logical definitions that can be used for logical inference (IHTSDO, 2015). SNOMED concepts are organized into a hierarchy of 'Is-a' relations. The top concept, **138875005 | SNOMED CT Concept (SNOMED RT+CTV3)** directly subsumes 19 high level concepts. This includes first order concepts such as **404684003 | Clinical finding (finding)**, and **123037004 | Body structure (body structure)**, which serve as the root of subhierarchies of concepts about entities directly relevant to and within the domain of healthcare. It includes also relations used amongst concepts in SNOMED CT as well as second order concepts that describe the structure of SNOMED CT rather than the structure of what the first-order concepts of SNOMED CT are about. Every SNOMED CT concept comes with *descriptions* one of which is selected as the *Fully Specified Name* (FSN) and which typically ends in a *semantic tag* (ST) that disambiguates it from other concepts that may have similar names (IHTSDO, 2015, p41). The ST also serves to indicate where the concept fits into the SNOMED CT concept hierarchy (IHTSDO, 2017). For example, the concepts [**35566002 | Hematoma (morphologic abnormality)**] and [**385494008 | Hematoma (disorder)**] have the STs 'morphologic abnormality' and 'disorder' attached to the *name* they have in common: *hematoma*. In the hierarchy, these concepts are ulti-

mately subsumed by the highest-level concepts for morphologic abnormalities and disorders respectively. Because STs are substrings added to *names* inside FSNs and are not represented separately as part of SNOMED CT's formal model, it is not easy to determine whether a tag on a concept should be taken to mean that the concept is *necessarily* part of the same sub-hierarchy as others with that tag. A concept's ST would strictly identify its place within the hierarchy if *each tag had a single, high-level corresponding concept that used it, and every concept using the tag was below that high-level concept in the hierarchy*. For instance: in the clinical finding hierarchy the highest finding concept, [**404684003 | Clinical finding (finding)**], subsumes all other findings.

The exact relationship between SNOMED CT's STs and concepts has thus far not been widely researched. In (Ceusters & Bona, 2016) we explored how the STs of concepts changed over time. We found in total 285 patterns according to which SNOMED CT concepts underwent changes in the STs assigned to them -- a change from no ST at all to a ST (43 patterns) counted also as a change. There were no patterns with more than 3 changes over time. Changes in STs were found to happen for a number of reasons. One is a change in SNOMED CT's concept model, for instance when distinctions are made that didn't exist in earlier versions, or different interpretations were introduced (e.g. the product / substance distinction). Such changes have a global impact on large parts of the ontology. Another reason is that concepts were in one or other way erroneous and had to be corrected. While doing these analyses, we were nevertheless hampered by the fact that the SNOMED CT documentation available from the IHTSDO webserver provides insufficient information on what the precise set of STs the SNOMED CT editors are working with might be. The information that a ST is that what appears at the end of a FSN between brackets (IHTSDO, 2015, p41) turned out not to be reliable. Historically, FSNs didn't have a ST at all as this was apparently introduced later as witnessed by the many changes in descriptions to that end. It was found that parsing anything that terminates a FSN between brackets leads to many false positives in older concepts, thus requiring manual inspection for disambiguation. The work presented here examines the January 31, 2017 International Release of SNOMED CT to investi-

* To whom correspondence should be addressed: jonathanbona@gmail.com

gate the extent to which SNOMED CT's use of STs is systematic and consistent with its placement of concepts that use those STs within the concept hierarchy. Research hypotheses driving this work are:

- (1) All STs are related to the concept system through a one-to-one correspondence between the ST and some high-level concept. Every concept that uses a particular ST t should be subsumed by that ST's 'corresponding concept' C_t , where C_t is the highest level concept that uses t . This hypothesis is motivated by the apparent change in terminology from 'semantic tag' in (IHTSDO, 2015) to 'hierarchy tag' in (IHTSDO, 2017).
- (2) We consider a concept to be 'mismatched' if it has the ST t but is not subsumed by the corresponding concept C_t .
- (3) Where such mismatches exist, they are due to errors in the concept's placement in the hierarchy or in its ST, and should be corrected in future releases.

This paper reports on techniques we have developed to detect mismatched concepts, categorize them, and extract patterns to understand how they change over time as new versions of SNOMED CT are released.

2 METHODS

We have developed computational procedures (1) to identify the concept that corresponds to an ST and (2) to facilitate answering questions about subsumption that involve considering all SNOMED CT concepts in each release. These are described in detail below.

2.1 Identifying tag corresponding concepts

In order to determine whether a concept C is mismatched it is necessary to know which concept is the corresponding concept for C 's ST. There does still not appear to be an official published mapping that lists the ST / concept correspondences for SNOMED CT. In many cases this correspondence may seem obvious to a human observer since for some tags there is a single high-level concept that uses the tag and whose name is the same as the tag. For example, one direct sub-concept of the top SNOMED CT Concept is **71388002 | Procedure (procedure)**. This concept is has the ST 'procedure' and its name in the FSN is the word 'Procedure'.

In other cases, the correspondence is less obvious. For instance, no direct sub-concept of the top concept is tagged 'morphologic abnormality', nor is there any concept whose name is exactly 'Morphologic abnormality'. The concept **118956008 | Body structure, altered from its original anatomical structure (morphologic abnormality)** is a child of **123037004 | Body structure (body structure)** and appears to be the highest concept (i.e. closest to the top) tagged with 'morphologic abnormality'.

We define therefore *the corresponding concept* for any ST t as: *the highest concept in the hierarchy that is tagged with t* . Note that this definition does not require tags to keep the same corresponding concept across releases.

Based on this we determine the corresponding concept C_t for each ST in a SNOMED release by:

- (1) Calculating the whole number *depth* for each concept C as the length of the shortest *Is-a* path from the top concept to C .
- (2) For each ST t , select from the set of concepts tagged with t the concept with the least *depth*, X_t .
- (3) Let $C_t = X_t$ if none of X_t 's ancestors is tagged with t . Otherwise let C_t be the ancestor of X_t that has the least *depth*.

Step 3 is necessary to handle special cases (**Fig. 1**) that arise from SNOMED CT's use of multiple inheritance caused by its *Is-a* hierarchy forming a directed acyclic graph with a single root node (*SNOMED CT Concept*) that has no edges coming into it (i.e. is not subsumed by any other concept). Such special cases occur whenever there is a concept with some ST t that is closest to the top as compared to all other concepts with ST t , and at the same time is also subsumed by another concept with ST t that has a longer shortest path to the top concept. Such patterns were found in some releases making it thus possible for a more general concept – i.e. higher up in the hierarchy – for a ST to be subsumed by less general concepts that use the same ST.

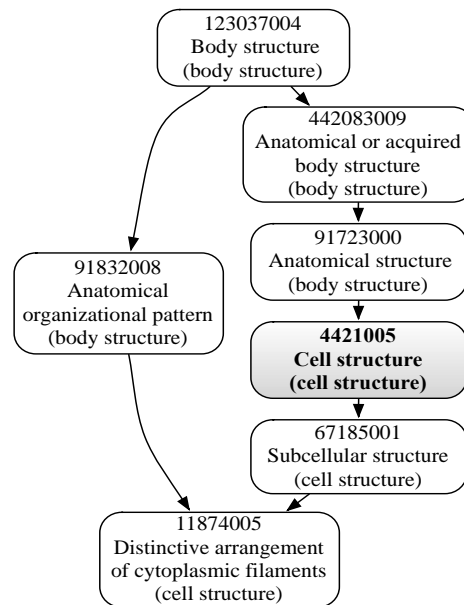


Fig. 1. Effect of concept multiple inheritance on ST hierarchy.

The output of this process is a mapping of STs to corresponding concepts for each release. This mapping is fairly stable across releases, though there are some changes, which we discuss more in the results section below.

2.2 Identifying mismatched concepts

Once a corresponding concept has been identified for each tag in each release, it is possible to find mismatched concepts by looking at each concept in turn to see whether it is subsumed by the corresponding concept for its ST. We developed computational procedures to do this.

In order to make use of the built-in subsumption reasoning provided by standard semantic web tools, we constructed an RDF/OWL model that represents SNOMED CT's concept hierarchy (300,000+ concepts connected by the *is-a* relation) and STs. Each concept is represented as an OWL class with separate annotations for its FSN and ST. Each of SNOMED's *Is-a* relations between concepts has a corresponding *rdf:subClassOf* assertion this representation. We built one such OWL file for each SNOMED CT release from January 2003 to January 2017. The identifiers (URIs) for each concept use a namespace that indicates the release version, e.g. `<http://ex.com/r20170131#64572001>` is an identifier for the concept with concept id 64572001 in the January 31 2017 release. These files were loaded into a single repository in a triple store database (Bishop et al., 2011) configured for RDFS+ inference that, upon loading, pre-computed subsumption for each hierarchy, resulting in a total of 185 million triples. This facilitates very fast retrieval of subsumption information using simple SPARQL queries, and allows us to instantly answer questions such as: *given a release R, a tag t, and a concept C, which concept - if any - are tagged with t in R, but not subsumed by C in R?* As an example, the following query retrieves the concept URI, label, and ST for every concept that is not subsumed by **64572001 | Disease (disorder)** even though it uses the corresponding tag.

```
PREFIX corr: <http://ex.com/r20170131#64572001>
PREFIX tagged: < http://ex.com/r20170131#tagged>
PREFIX : <http://ex.com/r20170131#>
SELECT ?conc ?l ?tag
WHERE {
  ?conc rdfs:label ?l .
  ?conc tagged: ?target_tag .
  corr: tagged: ?target_tag .
  ?conc tagged: ?tag .
  FILTER NOT EXISTS {?conc rdfs:subClassOf corr: }
}
```

Two flavors of mismatches were then looked for: (a) 'local' mismatching as defined in assumption (2) above which occurs within the scope of a specific release, and (b) 'global' mismatching in which the reference is the most recent release investigated. The group of globally mismatched concepts includes thus those concepts which have in at least one version *V* a semantic tag which is different from the one it has in the last version, whether or not it is locally mismatched in *V*.

2.3 Characterizing mismatched concepts

We then group locally mismatched concepts into categories based on the presence or absence of other mismatched con-

cepts among their subsumers. These were organized into a table with concepts as rows and SNOMED CT release dates as columns, with each cell indicating the concept's category for that release. We took of course into account that SNOMED CT concepts can be either active or inactive in a release and that a concept that is active in one release may be deactivated in the next one, for instance if the concept was deemed by SNOMED CT's editors to be no longer accurate or useful. Less commonly, a concept that is inactive at one release may be (re)activated at the next. We consider concepts to be not active in releases that precede their addition to SNOMED CT.

The categories into which concepts were classified were constructed by building up a three-character code ' _ _ _ ' where each character is a flag indicating whether a certain condition holds of the concept in that release. If a concept is inactive or did not yet exist at a release, then that concept was marked with the three-character empty code ' ' for that release. The following construction principles were used:

- The first character is 'Y' if the concept is subsumed by its ST's corresponding concept in this release (i.e. if it is NOT mismatched in the release), and 'N' otherwise.
- The second character is 'Y' if the concept has any ancestor concept that is NOT mismatched. It is 'N' if every ancestor of this concept is mismatched.
- The third character is 'Y' if the concept has any ancestor concept that IS mismatched. It is 'N' if no ancestor of this concept is mismatched.

Combinatorically, this would allow us to code for nine different situations including the inactive concepts. However, given the meanings assigned to these codes, some combinations are impossible. Ideally, every active concept in SNOMED would be in the 'YYN' category, indicating that the concept is properly matched to its ST's corresponding concept, as are all of the concepts above it. Possible codes for mismatched concepts are 'NYY' and 'NYN' while for non-mismatched concepts 'YYN' and 'YYY'. The latter indicates a concept that itself is not mismatched, but it is subsumed by at least one mismatched concept.

3 RESULTS

3.1 Corresponding concept mappings

We used the corresponding concept discovery procedure presented above to construct a table with [ST→concept] pairs and inspected this table manually to assess whether the mappings made sense. In the majority of cases, the ST turned out to be identical to the name of the corresponding concept modulo capitalization and spacing. Exceptions were: [SNOMED RT+CTV3→SNOMED CT Concept], [metadata → SNOMED CT Model Component], [Environment / location→Environment or geographical location], [Staging

	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
assessment scale	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
body structure	2	2	2	2	2	0	0	0	0	0	0	0	0	0	0
cell	25	4	4	4	4	4	1	0	0	0	0	0	0	0	0
cell structure	12	11	5	5	5	5	5	5	5	5	5	0	0	0	0
context-dependent category	856	870	686	567	-	-	-	-	-	-	-	-	-	-	-
disorder	1914	1477	2825	156	129	133	118	114	84	84	87	62	79	74	83
environment	63	63	63	33	0	0	0	0	0	0	0	0	0	0	0
ethnic group	2	1	0	0	0	0	0	0	0	0	0	0	0	0	0
event	8286	8307	8178	32	32	32	32	10	2	2	2	0	0	0	0
finding	2220	2022	2038	20	17	16	16	0	0	0	0	0	0	0	0
morphologic abnormality	13	14	5	5	5	5	5	4	4	4	4	4	0	0	0
navigational concept	553	479	453	307	280	260	241	3	3	2	1	1	1	1	0
observable entity	209	100	58	29	28	23	19	5	5	0	0	0	0	0	0
occupation	0	2	2	0	0	0	0	0	0	0	0	0	0	0	0
person	5	5	5	5	5	0	0	0	0	0	0	0	0	0	0
physical object	27	27	27	27	27	27	27	27	27	27	27	27	27	27	0
procedure	12	8	3	0	0	0	0	0	0	0	0	0	0	0	0
product	37	36	1	1	1	1	1	1	1	1	1	1	1	1	1
qualifier value	284	278	276	265	415	1	1	0	0	0	0	0	0	0	0
record artifact	-	-	-	2	1	0	0	0	0	0	0	0	0	0	0
regime/therapy	257	327	51	52	37	163	35	15	1	1	1	0	0	3	4
situation	-	-	-	-	540	478	464	175	1	1	1	1	1	1	0
special concept	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
specimen	2	2	2	2	2	0	0	0	0	0	0	0	0	0	0
substance	34	33	16	8	7	7	0	0	0	0	0	0	0	0	1
total mismatches	14814	14068	14700	1522	1537	1155	965	359	133	127	129	96	109	107	89
#mismatched tags	21	21	20	19	18	14	13	10	10	9	9	6	5	6	4

Table 1: Per-tag counts of globally mismatched concepts for the January release of each year

	0301	0307	0401	0407	0501	0507	0601	0607	0701	0707	0801	0807	0901	0907	1001
disorder	0	0	0	0	0	0	0	0	0	44	8	19	0	0	0
finding	4	1	0	5	0	0	0	0	0	0	0	0	0	0	0
observable entity	0	0	0	2	0	0	0	0	0	0	0	0	0	0	0
product	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1
regime/therapy	0	0	0	0	0	0	0	0	0	37	259	123	0	0	0
substance	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	1007	1101	1107	1201	1207	1301	1307	1401	1407	1501	1507	1601	1607	1701	Total
disorder	0	0	0	0	2	4	10	26	32	44	24	74	78	83	188
finding	0	0	0	0	0	0	0	0	0	0	0	0	0	0	10
observable entity	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
product	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
regime/therapy	0	0	0	0	0	0	0	0	0	0	3	3	3	4	263
substance	0	0	0	0	0	0	0	0	0	0	0	0	0	1	2

Table 2: Per-tag counts of locally mismatched concepts per release.

scale→Staging and scales], [situation→Situation with explicit context], [assessment scale→Assessment scales], [regime/therapy→Regimes and therapies], [cell→Entire cell], [morphologic abnormality→Body structure, altered from its original anatomical structure], [geographic location→Geographical and/or political region of the world], [product→Pharmaceutical / biologic product], and [disorder→Disease]. These seem to be plausible mappings. We then constructed a table with all tag corresponding concepts in every

release and examined this table for changes. We found that the majority of tag to corresponding concept pairings are stable over all releases. A few are absent initially but appeared when their tag was added to SNOMED, e.g. **415229000 | Racial group (racial group)** appears for the first time in January 2005. A few are present initially but disappeared when they were removed from SNOMED, e.g. **304813002 | Administrative values (administrative concept)**, which was removed as of the July 2010 release.

Some corresponding concepts had minor edits made to their FSNs, but we do not count this as a change in corresponding concept. Finally, one tag switched its corresponding concepts from one release to the next: the tag ‘finding’ initially had as its corresponding concept **246188002 | Finding (finding)** but this concept was deactivated in the January 2004 release and the ‘finding’ corresponding concept changed to **404684003 | Clinical finding (finding)**.

3.2 Mismatched concepts

After identifying all mismatched concepts for every ST in every release, we organized counts of mismatched concepts into two tables, one for global mismatching (**Table 1**) and one for local mismatching (**Table 2**), with one row per ST and one column per release (**Table 1**, for readability, contains only the counts for the January versions of each release). The number of globally mismatched concepts per release is generally decreasing over time, and has gone from 14,814 (5% of active concepts) in the January 2003 release to 89 (0.027%) in the January 2017 release. The number of global mismatches dropped dramatically from the July 2005 release (14,715 mismatches – 4.83% of active concepts) to the next release in January 2006 (1522 – 0.5%). This improvement is likely attributable to large changes in the hierarchy that involved changes in semantic tags for three hierarchies: ‘*disorder*’, ‘*event*’, and ‘*finding*’. This reorganization is documented in the SNOMED CT Editorial Guide’s section on Changes and historical notes: ‘In January 2006, a number of concepts from the | Clinical finding | *hierarchy* were moved to the *Event hierarchy*’ (IHTSDO, 2015, p294).

In the January 2017 release there are only four tags with mismatched concepts for a total of 89 mismatches: 83 are tagged ‘disorder’, four are tagged ‘regime/therapy’, one is tagged ‘product’, and one is tagged ‘substance’.

3.3 Mismatched disorders

Table 3 provides more detail about the categorization of the 188 locally mismatched ‘disorder’ concepts by release. This table was constructed by collecting all concepts that at least once in their lifetime were locally mismatched while having the semantic tag ‘disorder’.

The column marked ‘NE’ has for each release counts of the number of ‘disorder’ concepts that appear and are mismatched in some later SNOMED release, but that did not yet exist at the release for that row.

The ‘Inact.’ column counts how many were active concepts in an earlier release but were inactive at the row release.

The column ‘N | I’ is a sum of the previous two columns.

The ‘NYN’, ‘NYY’, ‘YYN’, and ‘YYY’ columns count the number of concepts that fall into each of those categories. In the January 2017 release, the 83 mismatched concepts fall into two of our categories: ‘NYN’ (69) and ‘NYY’ (14).

Release	NE	Inact.	N I	NYN	NYY	YYN	YYY	#Err.	%Err.	Total
20030131	78	0	78	0	0	110	0	0	0	110
20030731	78	0	78	0	0	110	0	0	0	110
20040131	75	0	75	0	0	113	0	0	0	113
20040731	75	0	75	0	0	113	0	0	0	113
20050131	75	0	75	0	0	113	0	0	0	113
20050731	75	0	75	0	0	113	0	0	0	113
20060131	75	0	75	0	0	113	0	0	0	113
20060731	75	0	75	0	0	113	0	0	0	113
20070131	73	0	73	0	0	115	0	0	0	115
20070731	66	0	66	18	26	78	0	44	36.07	122
20080131	61	0	61	6	2	119	0	8	6.3	127
20080731	60	0	60	1	18	109	0	19	14.84	128
20090131	60	0	60	0	0	128	0	0	0	128
20090731	60	0	60	0	0	128	0	0	0	128
20100131	60	15	75	0	0	113	0	0	0	113
20100731	59	15	74	0	0	114	0	0	0	114
20110131	59	15	74	0	0	114	0	0	0	114
20110731	59	15	74	0	0	114	0	0	0	114
20120131	59	15	74	0	0	114	0	0	0	114
20120731	58	15	73	2	0	113	0	2	1.74	115
20130131	56	15	71	4	0	113	0	4	3.42	117
20130731	50	15	65	10	0	113	0	10	8.13	123
20140131	40	15	55	17	9	107	0	26	19.55	133
20140731	35	15	50	22	10	106	0	32	23.19	138
20150131	21	15	36	34	10	108	0	44	28.95	152
20150731	12	16	28	24	0	136	0	24	15	160
20160131	6	16	22	65	9	92	0	74	44.58	166
20160731	5	16	21	67	11	89	0	78	46.71	167
20170131	0	16	16	69	14	89	0	83	48.26	172

Table 3. Locally mismatched ‘disorder’ concepts

4 DISCUSSION

Our hypothesis that SNOMED CT intends its STs to have a one-to-one correspondence between tags and certain high-level concepts is supported by:

- (1) the very existence of identifiable tag corresponding concepts (a single ‘highest’ concept for each tag that is close to the top concept and that in each case subsumes the vast majority of concepts that use the tag);
- (2) the generally very low occurrence of mismatched concepts.

Errors, however, remain present and sometimes are even introduced, interestingly as witnessed by our analysis of the locally mismatched disorder concepts (Table 2) increasingly more in more recent versions. A particularly illustrative example of this in the January 2017 release is the concept **109186003 | Sickle cell test kit (substance)** which turned out to be newly mismatched as there were no ‘substance’ concepts mismatched globally from 2009 until this release, and not even a locally mismatched ‘substance’ concept ever before. It is mismatched because it is not subsumed by the ‘substance’ tag’s corresponding concept **105590001 | Substance (substance)**. Indeed, the sickle cell test kit concept is directly

subsumed by **385387009 | Test kit (physical object)**, which has 29 other children that all have the words ‘test kit’ in their FSN and are correctly tagged with ‘physical object’ (e.g. **1109190001 | Virus test kit (physical object)**).

There are a number of ways for a mismatched concept to appear in a release. These include the addition of a new concept, re-activation of an old concept, and changes in the concept’s subsumption hierarchy. In the sickle cell test kit case changes in the hierarchy are responsible: in 2016 and earlier, the concept **385387009 | Test kit (physical object)** was itself mismatched, being subsumed by **105590001 | Substance (substance)** and not by the ‘physical object’ tag’s corresponding concept, **260787004 | Physical object (physical object)**. The test kit concept’s children were all tagged ‘substance’.

In 2017 the test kit concept was (correctly) moved to the physical object hierarchy, and it went from being mismatched to being not mismatched. 29 of its children had their FSNs changed to use the tag ‘physical object’. Though this move resolved one mismatch, another appeared: the sickle cell test kit concept became mismatched as a result, as shown in the two hierarchy excerpts in **Fig.4**. Most of the child concepts of **385387009 | Test kit (physical object)** are omitted here in the interest of space, as are child concepts of all the other concepts that appear here.

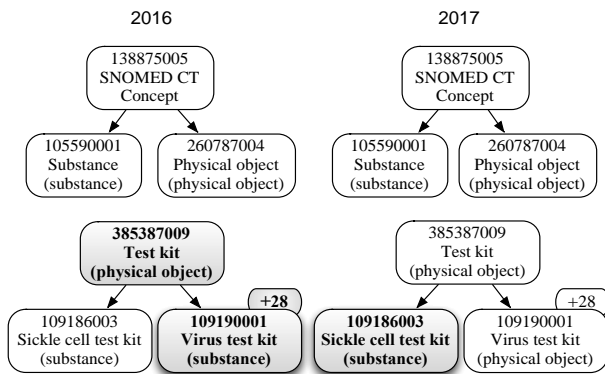


Fig.4. Test kit concept changes 2016 - 2017

5 CONCLUSION

We have successfully demonstrated that it is possible to implement an algorithm that maps semantic tags to corresponding SNOMED CT concepts. We applied this mapping in an analysis of all active concepts across SNOMED CT releases, assessing the extent to which the tags as used reflect the placement within the hierarchy of the concepts that use them both locally, and, in the spirit of Evolutionary Terminology Auditing (Ceusters 2009) with respect to the last version which functions as a gold standard. The results support our hypothesis that SNOMED CT indeed intends its semantic tags to have a one-to-one correspondence with certain high-

level concepts. The occurrence of mismatches between the semantic tags of lower-level concepts and their placement in another hierarchy than where expected according to the semantic tag is a sign that the SNOMED CT authoring tool is not equipped with a formal mechanism to keep the hierarchy consistent with the semantic tags. It is our recommendation that such mechanism would be implemented and the method developed here might be a good starting point in addition to other mechanisms for quality control that have been developed by third parties (Geller, Ochs, Perl, & Xu, 2012; Ochs et al., 2015).

ACKNOWLEDGEMENTS

This work was supported in part by Clinical and Translational Science Award NIH 1 UL1 TR001412-01 from the National Institutes of Health. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the NIDCR, the NLM or the National Institutes of Health.

REFERENCES

- Bishop, B., Kiryakov, A., Ognyanoff, D., Peikov, I., Tashev, Z., & Velkov, R. (2011). OWLIM: A family of scalable semantic repositories. *Semant. web*, 2(1), 33-42.
- Ceusters W. (2009) Applying Evolutionary Terminology Auditing to the Gene Ontology. *Journal of Biomedical Informatics* 42:518–529.
- Ceusters, W., & Bona, J. P. (2016). Analyzing SNOMED CT's Historical Data: Pitfalls and Possibilities. *AMIA Annu Symp Proc, 2016*, 361-370.
- Geller, J., Ochs, C., Perl, Y., & Xu, J. (2012). New abstraction networks and a new visualization tool in support of auditing the SNOMED CT content. *AMIA Annu Symp Proc, 2012*, 237-246.
- IHTSDO. (2015). *International Health Terminology Standards Development Organization - SNOMED CT® Technical Implementation Guide - January 2015 International Release (US English)*.
- IHTSDO. (2017). SNOMED CT Editorial Guide: 7.3.4. Hierarchy tag. Retrieved from <https://confluence.ihtsdotools.org/display/DOCEG/7.3.4+Hierarchy+tag>
- Ochs, C., Geller, J., Perl, Y., Chen, Y., Xu, J., Min, H., . . . Wei, Z. (2015). Scalable quality assurance for large SNOMED CT hierarchies using subject-based subtaxonomies. *J Am Med Inform Assoc*, 22(3), 507-518. doi:10.1136/amiajnl-2014-003151