Lexically Suggest, Logically Define: Quality Assurance of the Use of Qualifiers and Expected Results of Post-Coordination in SNOMED CT

Alan Rector & Luigi Iannone
School of Computer Science, University of Manchester, Manchester M13 9PL UK
rector@cs.manchester.ac.uk

Please cite as:
Rector, A., & Iannone, L. (2011). Lexically suggest, logically define: Quality assurance of the Use of qualifiers and expected results of post-coordination in SNOMED CT. *Journal of Biomedical Informatics* (in press, http://dx.doi.org/10.1016/j.jbi.2011.10.002)

(expanded version of AMIA 2011 distinguished paper award),
AMIA Abstract pg 1165 http://proceedings.amia.org/1lor3m/1lor3m/1.

Abstract

A study of the use of common qualifiers in SNOMED CT definitions and the resulting classification was undertaken using combined lexical and semantic techniques. The accuracy of SNOMED authors in formulating definitions for pre-coordinated concepts was taken as a proxy for the expected accuracy of users formulating post-coordinated expressions. The study focused on "acute" and "chronic" as used within a module based on the UMLS CORE Problem List and using the pattern of SNOMED CT's definition Acute disease and Chronic disease. Scripts were used to identify potential candidate concepts whose names suggested that they should be classified as acute or chronic findings. The potential candidates were filtered by local clinical experts to eliminate spurious lexical matches. Scripts were then use to determine which of the filtered candidates were not classified under acute or chronic findings as expected. The results were that 28% and 20% of candidate chronic and acute concepts, respectively, were not so classified. Of these candidate misclassifications, the large majority occurred because "acute" and "chronic" are sometimes specified by qualifiers for clinical course and sometimes for morphology, a fact mentioned but not fully detailed in the User Guide distributed with the SNOMED releases. This heterogeneous representation reflects a potential conflict between common usage in patient care and SNOMED's origins in pathology. Other incidental findings included questions about the qualifier hierarchies themselves and issues with the underlying model for anatomy. The effort required for the study was kept modest by using module extraction and scripts, showing that such quality assurance of SNOMED is practical. The results of a preliminary study using proxy measures must be taken with caution. However, the high rate of misclassification indicates that, until the specifications for qualifiers are better documented and/or brought more in line with common clinical usage, anyone attempting to use post-coordination in SNOMED CT must be aware that there are significant pitfalls.

Keywords: SNOMED, Terminology, Ontology, Knowledge Representation, Quality assurance

Highlights

- SNOMED expressions for acute and chronic findings were examined
- Lexical analyses of terms were compared with semantic analyses of definitions
  Acute/chronic were sometimes represented as clinical course, sometimes as morphology
- 25% contained irregularities that could cause errors in post-coordination
- However, only 5% contained no representation of acute/chronic

1 Introduction

Because it is implemented in a description logic, SNOMED CT [1] allows concepts either to be named and defined – *i.e.* "pre-coordinated" – or to be presented as an expression made up of other concepts – *i.e.* "post-coordinated". The underlying description logic determines the semantics and classification resulting from both pre- and post-coordination.

Using the description logic, it should make no difference, for example, whether a user expresses a concept using the pre-coordinated concept *Chronic pain* or by the equivalent post-coordinated expression, *Pain with clinical qualifier chronic* as shown, in simplified form, in Figure 1. The named concept and the post-coordinated expression should be found to be logically equivalent.

Pre-coordinated definition:

Chronic pain (finding) EquivalentTo:
Pain (finding) that *Clinical course (Attribute)* some *Chronic (qualifier value)*

Post-coordinated expression:

(*Pain (finding)* that *Clinical course (Attribute)* some *Chronic (qualifier value))*

Figure 1: Pre-coordinated definition and Post-coordinated expression for "chronic pain" (simplified)

SNOMED CT is increasingly being mandated in the US – *e.g.* for "meaningful use" [2] and for some purposes by the FDA [3, 4] – in the UK [5] and, to varying degrees, in other countries that have joined the maintaining organisation, the International Health Care Terminology Standards Development Organisation (for list see [6]).

There is increasing interest in and discussion of post-coordination, because it is impossible to enumerate all possible combinations of characteristics without causing a combinatorial explosion [7, 8]. For example, in SNOMED to define all possible concepts of the form "acute X" and "chronic X" would triple the number of concepts in SNOMED. For post-coordination to be used successfully, it must be assured that post-coordinated expressions will be classified correctly and matched to any pre-existing equivalent pre-coordinated concept.

However, the equivalence of post-coordinated and pre-coordinated expressions can fail for several reasons.

- The named expression – *e.g. Chronic pain* – may be specified wrongly or according to a different pattern than that used in the post-coordinated expression, *e.g.* it might use a different attribute or qualifier value. In this case the two are unlikely to be related at all by the logic.
- The named expression may be incompletely specified. For example, the qualifier "chronic" might have been omitted. In this case, the results are unpredictable. In some but not all cases, the named pre-coordinated class will appear as an ancestor of the post-coordinated expression.
- The named expression may be specified correctly, but only by a "partial definition" *i.e.* only by necessary conditions rather than by necessary and sufficient conditions. In natural language, this corresponds to defining *chronic pain* as *"a kind of pain that is chronic"* rather than *"any pain that is chronic"*. In SNOMED's formalism, nothing can be inferred to be a kind of, or equivalent to, a partially defined concept. Therefore, if the named pre-coordinated expression is only partially defined, it will at best be inferred to be a child of the post-coordinated expression. (This distinction is discussed in more detail in Section 2.1 and Figure 2.)

Others and ourselves elsewhere have analyzed different aspects of SNOMED – *e.g.* globally using structural methods [9, 10], for various ontological failings [11, 12], or for the accuracy of its inferred hierarchies in practical applications[13].

Our goal in this study was a preliminary experiment to estimate the likely consistency of post-coordination. Post-coordination is not yet sufficiently widely used to be able to experiment on a corpus of post-coordinated expressions directly. Therefore, as a proxy for the consistency of users in formulating post-coordinated expressions, we took the consistency of SNOMED authors in defining pre-coordinated concepts whose name included "acute" or "chronic." Although not all concepts of acute and chronic disease can be pre-coordinated without combinatorial explosion, many have been. Where they have been defined, we would expect them to be classified under the existing SNOMED concepts for *Acute* or *Chronic.* The exceptions are where the naming is in some way misleading or outdated as in "acute lymphocytic leukemia" which has come to indicate a specific morphology not reflected in the term.

To test the consistency of SNOMED authors in defining pre-coordinated concepts including "acute" or "chronic," the method developed by some of SNOMED CT's founders, "Lexically suggested logical closure" [14] seemed ideal, since users would be expected to be guided by the names when formulating post-coordinated expressions. (We rephrased the name slightly because "closure" has other meanings in the description logic community. Note also that technique was originally proposed for use with one of SNOMED CT's precursors, SNOMED RT [15], although the issues remain unchanged in SNOMED CT.) For a preliminary study, we restricted ourselves to a module based on the CORE Problem List Subset published by the UMLS [16] (http://www.nlm.nih.gov/research/umls/Snomed/core_subset.html).

Our primary finding is that, based on this lexical-semantic analysis, following the pattern used to define *Acute disease*[1] and *Chronic disease*[2] illustrated in abbreviated form in Figure 1 and shown in detail in Figure 4, we identified twenty-eight and twenty per cent misclassification for concepts whose names indicated "chronic" and "acute" respectively. A large fraction of these misclassifications related to SNOMED's handling of the intrinsic ambiguities in the notions of "acute" and "chronic." In the pattern illustrated in Figures 1 and 4, and the primary pattern illustrated in the SNOMED Users' Guide [17], "acute" and "chronic" are treated as "clinical courses". However, for acute and chronic inflammation and ulceration and in some other cases, "acute" and "chronic" are represented as morphology. This leads to the concepts for *Acute peptic ulcer*[3] and "*Acute colitis*" not being represented by the same pattern as "*Acute obstruction*"[4]. No simple post-coordinated expression or query – say for "Acute gastrointestinal disorder" – would find all three. There is mention of this issue in the User Guide, but no detailed specification of which pattern – *i.e.* attribute and qualifier - to use in which case. Furthermore, to add to the confusion, many conditions, *e.g. Acute bronchopneumonia*[5]*,* are defined using both qualifiers. (For details see 2.2 and 4.2.1.)

2 Background

2.1 Description logics

SNOMED CT is formulated in a description logic [18]. Description logics are subsets of first order logic specially designed to formulate definitions and to allow inference to be

[1] 2704003|Acute disease (disorder)|
[2] 27624003|Chronic disease (disorder)|
[3] 196682000|Acute peptic ulcer (disorder)|
[4] 197078006|Acute intestinal obstruction (disorder)|
[5] 123587001|Acute bronchopneumonia (disorder)|

computationally tractable. They form the basis of the Web Ontology Language (OWL) [19]. SNOMED is deliberately based on a relatively simple variant for which computation is guaranteed to be efficient, EL^{++} [20], which corresponds to the OWL-EL profile [21][6].

The fundamental principle of description logics is that one concept is a kind of another if, necessarily, all instances of the one is instances of the other. For example, any *Acute pneumonia* should be classified as a kind of *Pneumonia,* because any *Acute pneumonia* is necessarily a *Pneumonia.* Because all *Pneumonia* is a kind of *Disorder* it also is necessarily an *Acute disorder.* To this basic rule, SNOMED adds mechanisms to the effect that any disorder of the part is also a disorder of the whole. Hence, since the *Lung* – the site of *Pneumonia* – is a part of the *Respiratory system, Acute pneumonia* should also be classified as an *Acute respiratory disorder.*

A "reasoner" – sometimes called a "classifier" and really a special form of mechanical theorem prover – is used to classify the pre-coordinated concepts, *i.e.* to infer the poly-hierarchies of concepts based on their logical definitions. Used in this way, the reasoner can be thought of as a "terminology compiler" that takes the "stated form" of the definitions and compiles the hierarchies. The reasoner can also be used to determine where to place a post-coordinated expressions in the hierarchies and to determine if a post-coordinated expression is equivalent to any pre-coordinated concept. There are a number of reasoners available with varying characteristics and catering for varying subsets ("profiles") of description logics, *e.g.* $FaCT^{++}$ [22], Pellet [23], SNOROCKET [24], etc. Many use the standard OWL API [25], which allows them to be attached to various editors and software easily and provides a variety of other functions. (For a brief but more complete explanation of description logic semantics and notation see the on line Appendix V of [13] for a detailed account see [18]). SNOROCKET is particularly important for this discussion because it is specifically designed to be highly efficient for the EL^{++} subset used in SNOMED.

An important feature of description logics is that they allow concepts to be, in SNOMED's parlance, either fully or partially defined. To be fully defined means to be defined by necessary and sufficient conditions (EquivalentTo: in OWL); to be partially defined means to be defined by necessary conditions alone (SubClassOf: in OWL). Examples of complete and partial definitions for chronic pain (slightly simplified) are shown in Figure 2, which also serves to illustrate the correspondence between the OWL Manchester syntax used in this paper, the variable free description logic syntax used in many description logic references, and standard first order logic notation.

[6] In fact, SNOMED's formalism is slightly less expressive than OWL-EL, because it does not use disjointness axioms.

OWL (Manchester Syntax):

Partial definition

Chronic pain (finding) EquivalentTo:
Pain (finding) that *Clinical course (Attribute)* some *Chronic (qualifier value)*

Complete definition

Chronic pain (finding) SubClassOf:
Pain (finding) that *Clinical course (Attribute)* some *Chronic (qualifier value)*

Description Logic variable free syntax:

Partial definition

Chronic_pain ⊑ *Pain* ⊓ ∃ qualified_by.*Chronic*

Complete definition

Chronic*_pain* ≡ *Pain* ⊓ ∃ *qualified_by.Chronic*,

Standard First order Logic syntax:

Partial definition

∀ x . *Chronic_pain*(x) → *Pain*(x) ∧ ∃ y. *Chronic*(y) ∧ *qualified_by*(x, y)

Complete definition

∀ x . *Chronic_pain*(x) ←→ *Pain*(x) ∧ ∃ y. Chronic(y) ∧ *qualified_by*(x, y)

Figure 2: Partial and complete definitions of chronic pain in OWL and in two other syntaxes

2.2 Note on notation and typography

Throughout this paper we shall use the OWL form of SNOMED, as converted in the Perl script provided with the distribution, and the Manchester OWL Syntax [26]. To improve readability, SNOMED terms will be printed in italics rather than placed in single quotes. To avoid ambiguity and maintain consistency with what is shown in the OWL conversion of SNOMED, SNOMED fully specified names will be used in all figures. Full SNOMED identifiers are given in footnotes for concepts appearing in the text.

2.3 The usual SNOMED patterns for definitions as expressed in OWL

Historically, SNOMED CT evolved from SNOMED International (and still earlier versions) in which there were four original axes – topography (anatomy), morphology, etiology, and function, to which numerous other were added over the years [27]. In SNOMED CT the most important of these for this paper manifest as *Site*[7] (anatomy), *Morphology*[8], and qualifiers, which include the qualifier *Clinical course*[9], which is the main qualifier for *Acute*[10] and *Chronic*[11]. For example, the definition of Atelectasis[12] is shown in Figure 3. This means "Atelectasis is a kind of Disorder with the morphology collapse and the finding site of lung". Acute atelectasis is defined by adding the qualifier acute ("Sudden onset AND/or short duration"[13]), in a separate definition as shown. This means: "Acute atelectasis is any Atelectasis that has the clinical course of Sudden onset AND/OR acute duration"). The

[7] 363698007|Finding site (attribute)|
[8] 116676008|Associated morphology (attribute)|
[9] 263502005|Cinical course (attribute)|
[10] 424124008|Sudden onset AND/OR short duration (qualifier value)|
[11] 90734009|Chronic (qualifier value)|
[12] 46621007|Atelectasis (disorder)|
[13] 424124008|Sudden onset AND/OR short duration (qualifier value)|

definition of Atelectasis is partial defined because it uses SubClassOf; the definition of Acute atelectasis is a fully defined because it uses EquivalentTo (even though one of the concepts in the definition is only partially defined).

The property *RoleGroup* serves to group together closely related characteristics[28]. It must be included in all concept definitions for consistency with the minority of concepts in which it makes a difference. However, none of the examples discussed in this paper depend on the use of *RoleGroups*. Furthermore, their use and meaning remains controversial and may be subject to change [29]. For purposes of understanding this paper, they can be treated as technical artifacts and disregarded, although anyone wishing to replicate or extend these experiments will need to include them.

Atelectasis (disorder) SubClassOf:
Disorder of lung (disorder) and
(*RoleGroup* some
((*Associated morphology (attribute)* some *Collapse (morphologic abnormality)*) and
(*Finding site (attribute)* some *Lung structure (body structure)*)))

Acute atelectasis (disorder) EquivalentTo:
Atelectasis (disorder) and
(*RoleGroup* some
(*Clinical course (attribute)* some *Sudden onset AND/OR short duration (qualifier value)*))

Figure 3: SNOMED definition of Ateletasis – a morphology combined with a site – and Acute Atelectasis – Atelectasis with the added modifier for acute.

2.4 The "stated form" and classification in SNOMED

Using a description logic means that SNOMED is authored in two steps.

- First, the authors formulate the logical definitions in the description logic or "stated form."
- Second, a reasoner or classifier is used to infer the concept hierarchy from the definitions and descriptions. It is the inferred hierarchy that forms the basis of the distribution files that are seen by users.

Using a description logic has many advantages. It is the key to post-coordination and to expressing the relation between findings and their anatomical sites. It means that if an error is found, for example in the anatomy section, a single change may correct the classification of dozens or even hundreds of concepts.

However, using a description logic brings with it at least three difficulties:

- The phrasing of the fully specified name for a concept may not correspond to the definition in the description logic.
- The full consequences of logical definitions to classification by a reasoner may not be obvious to authors, any more than the full consequences of statements in a programming language to the behaviour of the compiled program are always obvious to programmers. Unexpected "bugs" can occur in either.
- Errors must be corrected in the "stated form," which is analogous to a program's source code. Trying to correct errors directly in the inferred or distributed form risks having them reappear the next time the classifier is run. It is as futile as trying to fix a program by patching the output of the compiler – a practice once common but long since abandoned as being unmaintainable.

2.5 Representation of "Acute" and "Chronic" in SNOMED CT

The existing definition for *Acute disease*[14] and *Chronic disease*[15] are shown in the first part of Figure 4. This is the primary representation described in the SNOMED User Guide[30], although it includes the addendum:

> *The word acute has more than one meaning, and the meanings are often overlapping or unclear. The word acute may imply rapid onset, short duration, or high severity; in some circumstances it might be used to mean all of these. For morphological terms it may also imply the kind of morphology associated with the speed of onset. |Acute inflammation (morphologic abnormality)| does not necessarily have CLINICAL COURSE | Sudden onset AND/OR short duration |, but rather implies polymorphonuclear infiltration; likewise |Chronic inflammation (morphologic abnormality)| implies mononuclear cell infiltration, not necessarily a chronic course, although inflammation with a chronic course is highly correlated with a lymphocytic infiltration. (Ibid. pg 32).*

However, no detailed description of what counts as acute or chronic inflammation morphology is given nor are there any guidelines for when *Clinical course*[16], or *Morphology*[17] or both should be used. For example, the fact that acute and chronic ulceration are kinds of acute and chronic inflammation morphology, respectively, is not detailed, The information is sufficiently vague that we have assumed that when coding notions such as acute or chronic ulcer, pneumonia, obstruction, bronchitis, etc. most users would follow the pattern using the *Clinical course* attribute as used in the definition of *Acute disease* and *Chronic disease*.

Existing definitions for chronic and acute disease

Chronic disease (disorder) EquivalentTo:
    Disease (disorder) that
        (*RoleGroup* some (*Clinical course (attribute)* some *Chronic (qualifier value)*))

Chronic clinical finding (finding) EquivalentTo:
    Clinical finding (finding) that
        (*RoleGroup* some (*Clinical course (attribute)* some *Chronic (qualifier value)*))

Broadened definitions for chronic and acute clinical finding

Acute clinical finding (finding) EquivalentTo:
    Clinical finding (finding) that
        (*RoleGroup* some
            (*Clinical course (attribute)* some *Sudden onset AND/OR short duration (qualifier value)*))

Chronic clinical finding (finding) EquivalentTo:
    Clinical finding (finding) that
        (*RoleGroup* some
            (*Clinical course (attribute)* some *Chronic (qualifier value)*))

Figure 4: Definitions of acute and chronic disease and extension to clinical finding

3 Materials

Several recent developments have made it much easier to study and manipulate SNOMED.

- *The IHTSDO has released the stated form and a Perl script to convert it into OWL syntax*, which makes it possible to manipulate it using standard tools for OWL and other description logics.

[14] 2704003|Acute disease(disorder)|
[15] 27624003|Chronic disease (disorder)|
[16] 263502005|Cinical course (attribute)|
[17] 116676008|Associated morphology (attribute)|

- *The UMLS now maintains a CORE Problem List Subset* of the roughly 8500 most commonly used codes in several major US hospitals [16].
- *Methods to extract "modules" from description logic models based on "signatures"*. A "signature" is just a set of concepts. A "module" is a subset of the statements (axioms) in the model sufficient to guarantee that all of the inferences that would have been made amongst the concepts in the signature using the entire model will be also be inferred in the extracted module [31]. Most analyses can, therefore, be carried out on the extracted module, which is usually one or two orders of magnitude smaller than the full model.
- *A scripting language, OPPL* [32] has been developed that allows lexical and semantic criteria to be combined so as to be able to implement the "Lexically suggest, logically define" strategy interactively and to make it replicable.
- *Techniques for comparing and patching OWL models*, so that the changes ("diff") between two versions can be expressed as a "patch" that can be applied to the same or other versions.
- *Faster reasoners and classifiers and more powerful machines* that make it practical to work interactively with modules extracted from SNOMED and even with SNOMED as a whole with delays of seconds or minutes rather than hours or days. (For a review of reasoners with respect to SNOMED see [33].)

Taking advantage of these developments:

For the SNOMED stated form, we used the July 31 2010 IHTSDO (international) release of SNOMED CT converted to OWL using the Perl script provided with the release.

For the UMLS CORE Subset [7], we used the August 2010 release. (http://www.nlm.nih.gov/research/umls/Snomed/core_subset.html). (Note that there appears to be a few concepts in the subset missing from the SNOMED release used, but this does not affect the study results.)

For module extraction, we used the methods in the OWL API [25] within a simple package developed locally and made publicly available on the web. (http://owl.cs.manchester.ac.uk/snomed**).**

For browsing and editing the resulting OWL files, we used Protégé 4.1 (http://protege.stanford.org)

For classification, we used the Pellet 2 and FaCT++ reasoners integrated with Protégé 4.1. (Earlier studies used SNOROCKET [24] (http://aehrc.com/hie/snorocket.html), which is significantly faster but as of the time of writing worked only with the older Protégé 4.0, which did not support OPPL2.)

For scripting, we used the latest release of OPPL2 integrated into the Protégé 4.1 environment [14] (http://oppl2.sourceforge.net).

For validation of issues against the complete SNOMED release, we used primarily the SNOB (http://snob.eggbird.eu/) browser because it gives the easiest views upwards in the hierarchies.

For comparison and patching, we used the locally developed OWLPatch package, which is publicly available from Manchester [34] (http://owl.cs.manchester.ac.uk/patch/).

For computation, all analyses were run on a 2.53 GHz MacPro with 8GB of RAM.

For clinical expertise, advice was sought informally from clinical colleagues on collaborating projects, but no formal studies of clinical consensus were performed at this stage.

4 Methods

4.1 Overall approach

We used an adaptation of Campbell *et al*.'s [6] "Lexically suggested logical closure methods". To broaden the study slightly, we first defined concepts for *Chronic clinical finding*" and "*Acute clinical finding*" by substituting Clinical finding in SNOMED's definition of *Acute disease*[18] and "*Chronic disease*[19] as shown in the second part of Figure 4. (In SNOMED *Disease*[20] is a kind of *Clinical finding*[21]*.)*

Our assumption was that for anyone but the most sophisticated SNOMED expert, this would be the pattern that they would follow when forming post-coordinated expressions for acute or chronic diseases or findings. We then examined how many of the existing SNOMED pre-coordinated concepts whose fully specified names suggested that they were chronic or acute were not classified by the description logic classifier under *Acute* or *Chronic finding*. The methods were developed using "Chronic" and then validated by applying them to "Acute". In the description that follows, we focus on "Chronic" for brevity. A table summarizing the results for both "Chronic" and "Acute" is given in the Results section (Section 5).

The preliminary step was to extract a "module," as described in Section 3 above, from the SNOMED stated form using as a "signature" the UMLS CORE Problem List Subset augmented by fully defined classes for *Chronic clinical finding* and *Acute clinical finding*, as defined in the second part of Figure 4.

The overall strategy then was to:

- Identify potential "candidate" misclassified concepts whose names suggested that they should have been classified under *Acute* or *Chronic finding* but whose logical definitions did not lead to them being so classified by the reasoner.
- Filter out spurious potential candidates with the help of collaborating clinicians where, despite the name, they did not consider it to be appropriate to classify the potential candidates under *Acute* or *Chronic* finding.
- Analyse manually the reasons for the misclassification for each of the remaining candidates, and then prove the analysis by making the changes indicated by the analysis and testing that the candidate was then classified as expected under *Acute* or *Chronic finding.*

The number of candidates was recorded at each stage, both in absolute numbers and as a percentage of the total potential candidates identified. These figures are reported in the Results section. Following the extraction of the module, the analysis proceeded in the following steps.

0. Counting all concepts whose fully specified names began with "Acute" or "Chronic" to establish a base line.
1. Initial lexical/semantic search for potential candidates whose preferred term begins with "Acute" or "Chronic" but which were not classified under *Acute* or *Chronic finding* using the OPPL scripting language.
2. Exclusion of spurious potential candidates in collaboration with clinicians.
3. Examination of remaining candidates to identify systematic issues and creation of axioms to deal with these issues.

[18] 2704003|Acute disease(disorder)|
[19] 27624003|Chronic disease (disorder)|
[20] 64572001|Disease (disorder)|
[21] 404684003|Clinical finding (finding)|

4. Manual modifications of the remaining candidates to deal with incomplete or incorrect representations.
5. Repetition of Steps 1-4 for concepts whose names contained, rather than began with, "acute" or "chronic."
6. Conversion of the definitions of as many as possible of the top-level candidate from partial to complete, so as to reduce duplication.
7. Application of all of the above changes to the complete SNOMED stated form, and identification the number of potential candidates remaining in order to get an estimate of the percentage of the analogous task on all of SNOMED completed by working only on the module.

4.2 Step by step procedures

4.2.1 Step 1: Initial lexical/semantic search for candidates whose preferred term begins with "Acute" or "Chronic" but which were not classified under *Acute* or *Chronic finding*

```
?C:CLASS=MATCH("'Chronic.*")
SELECT ?C SubClassOf 'Clinical finding (finding)'
WHERE FAIL ?C SubClassOf  'Chronic clinical finding (finding)'
BEGIN
ADD ?C SubClassOf Candidate
END;
```

Figure 5a. Script to find all concepts with names beginning "Chronic." not classified under *Chronic findings* and make them subclasses of *Candidate*.

A preliminary search was performed for concepts whose names began with "Chronic" but were not classified as *Chronic clinical finding* using the OPPL script shown In Figure 5a. The regular expression "'Chronic.*" is used for the lexical part of the search and the SELECT statement for the semantic part of the search. The number of candidates identified is reported by OPPL and appears in Table 1 Step 1.

Note that the SELECT statement refers to the inferred hierarchy and SubClassOf finds all descendants not just the immediate children. Note also that FAIL in OPPL means "could not be found" as in SQL (*i.e.* "negation as failure"). By contrast, not in OWL means "provably false", *i.e.* impossible.

The resulting ontology was then classified so that the candidates were organized into inferred subtrees as shown in Figure 5b. The number of subtrees is shown in brackets in Table 1 and is a measure of the number of modifications required to correct the misclassification, since modifications are inherited by descendants. (For details of modifications, see steps 3-4.)

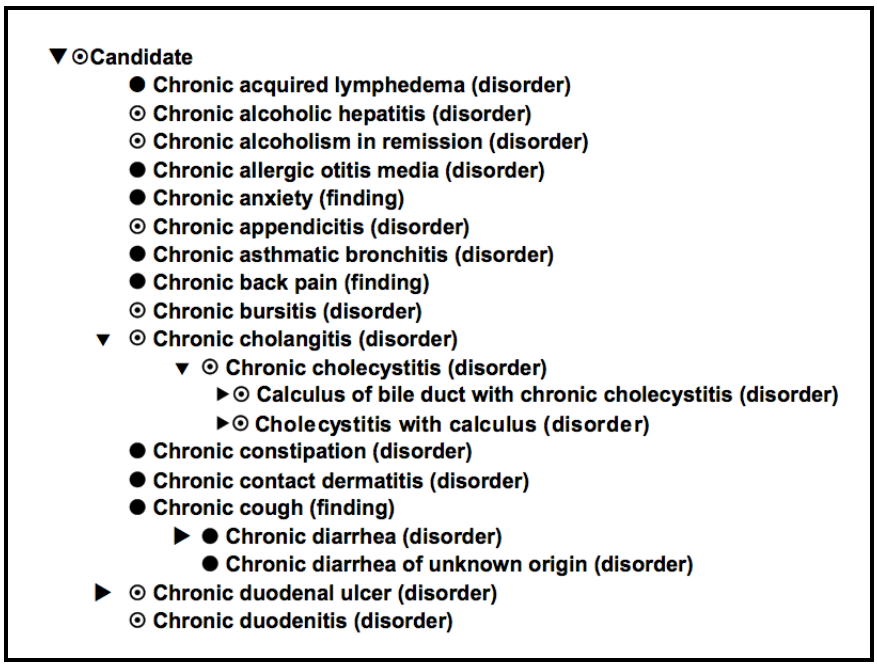

Figure 5b: Portion of classified list of candidates from search in Figure 5a. Solid dots indicate partially defined concepts; open dots indicate fully defined concepts. (For original screen shot, see on-line appendix.)

After classification, each candidate was examined manually. This identified three cases:

i. Cases where the words *Chronic* or *Acute* have clearly come to have a morphological (*i.e.* histological) meaning e.g. *Chronic myeloid leukemia*[22].

ii. Cases where *Chronic* or *Acute* are represented by *morphology* rather than the *course* qualifiers, but where informants questioned that this was appropriate for patient care, *e.g. Chronic peptic ulcer*[23].

iii. Cases where the representation was incomplete where there was no mention of *chronic* or *acute* in any qualifier, *e.g. Chronic back pain*[24].

4.2.2 Step 2: Exclusion of spurious candidates in collaboration with clinicians:

Concepts judged definitely to be spurious included all and only leukemias and myeloproliferative disorders. The script was therefore modified to exclude these conditions as shown in Figure 6.

```
?C:CLASS=MATCH("'Chronic.*")
SELECT ?C SubClassOf 'Clinical finding (finding)'
WHERE FAIL ?C SubClassOf  'Chronic clinical finding (finding)',
          FAIL ?C SubClassOf 'Myeloproliferative disorder (disorder)',
          FAIL ?C SubClassOf 'Leukemia, disease (disorder)'
BEGIN
ADD ?C SubClassOf Candidate
END;
```

Figure 6: Revised script to exclude leukemias and myeloproliferative disorders.

[22] 92818009|Chronic myeloid leukemia, disease (disorder)|
[23] 128287004|Chronic peptic ulcer (disorder)|
[24] 134407002|Chronic back pain (finding)|

4.2.3 Step 3: Creation of axioms to deal with case ii)

For the concepts that fell into case ii), we examined each *morphology* that appeared to imply a *clinical course*. For example, *Chronic* in *ulcer disease* is defined by the having the morphology *Chronic ulcer*[25], whereas *Chronic liver disease* is defined as having clinical course *Chronic*. There seemed no reason to our collaborators to treat *Chronic peptic ulcer* differently from *Chronic liver disease*[26], nor why a different pattern of post-coordination should be required to represent the one rather than the other. Figure 7a shows a examples of such concepts.

In order to assess the extent of this phenomenon, we created axioms to the effect that all conditions with these morphologies had a corresponding course, *e.g. Chronic inflammatory morphology* also had a *chronic course*. For experimental purposes, we did so even though this is explicitly contrary to SNOMED's *Style guide: Clinical Findings* [35] (which is different from the more generic *User Guide* [8] distributed with the release). The list of *morphologies* considered in this way is shown in Figure 7b. Of these, the morphology about which there is the most question is *Acute and chronic inflammation* because it includes *subacute* conditions. However, there is no separate qualifier for subacute, so that to treat them differently would require modifying the morphologic abnormality hierarchy. (Note that *leukemias* and *myeloproliferative* disorders had been excluded in step 2.0.)

196652006	|Acute duodenal ulcer (disorder)|
91357005	|Acute endocarditis (disorder)|
67602004	|Acute peritonitis (disorder)|
128286008	|Chronic duodenal ulcer (disorder)|
4851007	|Chronic endocarditis (disorder)|
87510000	|Chronic peritonitis (disorder)|

Figure 7a Example concepts where acute or chronic is represented by morphology rather than clinical course and are therefore not found by the standard pattern of post-coordination**.**

409777003	|Chronic inflammatory morphology (morphologic abnormality)|
405719001	|Chronic ulcer (morphologic abnormality)|
75889009	|Acute and chronic inflammation (morphologic abnormality)|
4532008	|Acute inflammation (morphologic abnormality)|
409776007	|Acute inflammatory morphology (morphologic abnormality)|
2052000	|Acute necrosis (morphologic abnormality)|
26317001	|Acute ulcer (morphologic abnormality)|

Figure 7b Morphologies hypothesized to imply chronic or acute course and used in axioms analogous to that in **Figure 8**.

Axiom_chronic_inflammatory_morphology_implies_chronic
EquivalentTo:
Clinical finding (finding) and
(*RoleGroup* some
(*Associated morphology (attribute)* some *Chronic inflammatory morphology (morphologic abnormality)*))
SubClassOf:
RoleGroup some (*Clinical course (attribute)* some *Chronic (qualifier value)*)

Figure 8: Experimental axiom that *chronic inflammatory morphology* implies *chronic course* used to test the extent of the effect of the use of morphology axioms. (For original Protégé screen-shot, see on-line appendix.)

[25] 405719001|Chronic ulcer (morphologic abnormality)|
[26] 328383001|Chronic liver disease (disorder)|

There are two ways to formulate axioms in OWL. They can be anonymous or they can be named defined classes that have further subclass or equivalent class axioms. The second type may be paraphrased: "Anything that satisfies the definition of this class must also have this additional qualifier." For this study, we chose the second form because it is easy to see which concepts have been directly affected by each axiom since they appear under that axiom in the inferred hierarchy. An example axiom is shown in Figure 8 in Manchester syntax and in the online appendix as it appears in Protégé.

4.2.4 Step 4: Manual modification to deal with case iii

Original under-specified partial definition:
Chronic anxiety state (finding) SubClassOf:
Anxiety state (finding)

Modified by adding qualifier for chronic to create a more complete partial definition:
Chronic anxiety state (finding) SubClassOf:
Anxiety state (finding) and
(*RoleGroup some (Clinical course (attribute) some Chronic (qualifier value)*)

Converted to a complete definition for a fully defined concept:
Chronic anxiety state (finding) EquivalentTo:
Anxiety state (finding) and
(*RoleGroup* some *(Clinical course (attribute) some Chronic (qualifier value)*))

Figure 9: Simple addition of the clinical course qualifier to create a better specified partially defined concept, followed by conversion to a fully defined concept.

The remaining concepts involve underspecified or incorrect definitions – case iii – and must be dealt with manually as shown in Figures 9. The list of candidate axioms was first cleared and then recalculated on the ontology including the axioms added in Step 3 using OPPL scripts. In each case – chronic and acute – fewer than twenty concepts remained to be dealt with.

The appropriate use of *RoleGroups* in some situations was unclear (See discussion in Section 2.5). For this study, all cases the correct qualifier was simply added nested within its own *RoleGroup* for consistency with usage with other definitions in SNOMED. Without the nesting within a RoleGroup, the reasoner does not infer the intended classification. (Where *RoleGroups* already existed, the alternative of nesting the qualifier within the existing *RoleGroup* was tried, but in no case found to result in a different classification.)

If the base definition was sufficient, the concept was then converted from partially to fully defined as shown. If not, the definition was left partial, pending more precise definition of the base concept.

4.2.5 Step 5: Repeat steps 2-4 for Clinical findings whose name contain the word "chronic" or "acute"

Steps 2-4 were repeated for concepts with names containing, rather than beginning with, "chronic" or "acute". Rather than describing these steps again, we simply label them steps 5a-5c in Table 1. There are many fewer candidates at this stage but a higher proportion of false positives.

4.2.6 Step 6: Convert as many as possible partial to complete definitions

The final step was to convert as many as possible of the of the top-level candidate classes (as shown for Chronic in Figure 5b) to defined classes. This is, inevitably, a manual procedure because it requires judgment as to whether or not the class is adequately defined by the existing expression.

4.2.7 Step 7: Assess number of remaining candidates in complete SNOMED

The OWLPatch package tools owldiff and owlpatch were used to extract the changes from the module and apply them to the entire SNOMED stated form. The OPPL scripts were then used to determine the number of remaining candidates in the stated form as a whole, excluding the previously identified spurious matches for *leukemias* and *myeloproliferative disorders* plus the further spurious match to *non-human disorders*. The number of remaining concepts to be dealt with was recorded as an estimate of the percentage of the total number of candidates in SNOMED as a whole that had been affected by dealing only with the those in extracted module. Since many of the changes were generic or affected high-level concepts with many descendants, it was expected that dealing with only the module would address disproportionately many candidates in SNOMED as a whole.

5 Results

The initial module extracted contained 33,000 concepts or just less than ten per cent of the total SNOMED CT corpus.

5.1 Step by step results of remainder of the analysis

The numbers of candidate concepts found by the lexical search and remaining after each step is shown in Table 1. The raw numbers of candidates at each stage are given along with the number of subtrees into which they were grouped by the classifier. The table is in two parts.

- The first records the results of the four basic steps for candidates with names beginning with the qualifier "Chronic" or "Acute". (Steps 1-4)
- The second records the additional effect if extending the lexical pattern to cover name containing, rather than beginning with, the keywords. (Step 5)

As can be seen, 30% and 22% of the candidates matching the lexical criteria were not found to match the semantic criteria of being subsumed by *Chronic finding* or *Acute finding*, respectively (Step 1). These numbers drop slightly to 28% and 20% when spurious matches are excluded (Step 2). Of the remainder, the majority relate to the issues between morphology and course (Step 3). The remaining 5% and 3% for "chronic" and "acute" respectively represent errors and omissions that must be dealt with manually (Step 4).

Table 2 summarises the effect of applying the changes made in the modules to SNOMED as a whole (Step 7). Slightly more new candidates remain after applying the changes than were in the module, suggesting that the corrections made in the module covered 41% and 44% of the similar cases in SNOMED as a whole.

A few other results are noteworthy. All candidates that were excluded as spurious when matching for the occurrence of the qualifier anywhere in the name (Step 5c) were sequelle or complications of chronic or acute diseases, and so included the qualifier in their name without themselves being acute or chronic. The largest group of valid new candidates in Step 5c fell under the *Acute and chronic inflammation*[27] and so were recognized by Steps 1-4 for *Acute* but not for *Chronic*. Of the remainder, the majority began with words like "Primary" or "Idiopathic."

[27] 75889009|Acute and chronic inflammation (morphologic abnormality)|

After Step Number	N % (Number of subtrees)		Comments
	Chronic	Acute	
Step 0: Total number of concepts with names beginning with Chronic/Acute	368 100% (64 subtrees))	450 100% (55 subtrees)	Baseline
Step 1: Initial candidates:	110 30% (62 subtrees)	101 22% (54 subtrees)	Concepts from Step 0 not classified under 'Chronic finding'.
Step 2: After exclusion of spurious candidates	103 28% (59 subtrees)	92 20% (52 subtrees)	Excludes leukemias and myeloproliferative disorders
Step 3 After addition of axioms	18 5% (16 subtrees)	14 3% (12 subtrees)	Remaining to be dealt with manually
Step 4: After manual changes to remaining candidates	0 (0 subtrees)	0 (0 subtrees)	All changes made; checked by script
Step 5a: Analysis of remaining names containing 'chronic" or "acute" anywhere: Initial candidates	21 100% (17 subtrees)	1 (1 subtrees)	Baseline for phase 2
Step 5b: After exclusion of spurious candidates	15 71% (11 subtrees)	0 (0 subtrees)	See discussion
Step 5c: After manual changes to remaining candidates	0 (0 subtrees)	0 (0 subtrees)	All changes made; checked by script

Table 1: Results of Steps 1-5, primary analysis (Step 6 does not affect figures)

Action	Chronic	Acute	Comments
1 Remaining candidates in complete SNOMED after applying changes from repaired module	146	117	Excludes leukemias, myeloproliferative and non-human diseases
2 Estimate of total number of candidates in entire corpus	249	209	Step 2 (Table 1) + Step 7 (Row 1 Table 2)
3 Estimate of the percentage of candidate errors in the entire SNOMED corpus dealt with by correcting the module	41%	44%	1 – (Row 1 / Row 2)

Table 2: Results of Step 7: application of changes to entire SNOMED corpus

5.2 Other issues raised

There are a number of other issues raised. The most important concern the qualifier hierarchy for *clinical course* shown in Figure 10. In particular, collaborating clinicians questioned whether *Intermittent*, *Seasonal*, or *Recurrent* should necessarily be classified under Chronic. This requires further investigation combining linguistic, ontological, and clinical expertise. We did not have the resources of a detailed study but took preliminary soundings. Our limited range of informants assented, for example, to both "intermittent acute pain" and "chronic intermittent pain", and even to "recurrent intermittent acute pain'. However, they dissented from "intermittent chronic pain" and questioned "acute intermittent pain." If confirmed, this

suggests that not all *intermittent* findings are *chronic*, and that the qualifier *intermittent* should not, therefore, be a kind of the qualifier for *chronic*.

Amongst the further issues that require further study are:

- The definition of *Chronic ulcer morphology*[28] includes *Pressure ulcer morphology*[29] which in turn includes *Decubitus ulcer morphology*[30]. All *pressure* and *decubitus* ulcers are, therefore, classified as *chronic*. Some informants questioned this. It appears to be part of a more generic problem that many qualifiers that are treated as morphology might seem to be more naturally treated as etiology. This probably stems from SNOMED's legacy in pathology system.
- *Chronic ulcerative inflammation morphology*[31] is neither a ancestor nor descendant of *Chronic ulcer morphology*,[32] and *Chronic ulcer morphology* is not a kind of *Chronic inflammation morphology*[33] – nor of any other morphology associated with the notion "chronic." This suggests an error in the qualifier hierarchy.
- The *Lumbar region back*[34] is defined to be part of the *Abdominal wall*[35], which is defined to be part of the *Abdomen*. The result is that *Low back pain* is classified as a kind of *Abdominal pain*[36]. This is an example of a wider problem in SNOMED that conflates bounds with parts. Repair awaits wider restructuring of SNOMED's anatomy, which is known to contain other such systematic errors [13] and is said to be under review by the IHTSDO. (Personal communication, Kent Spackman, October, 2010)

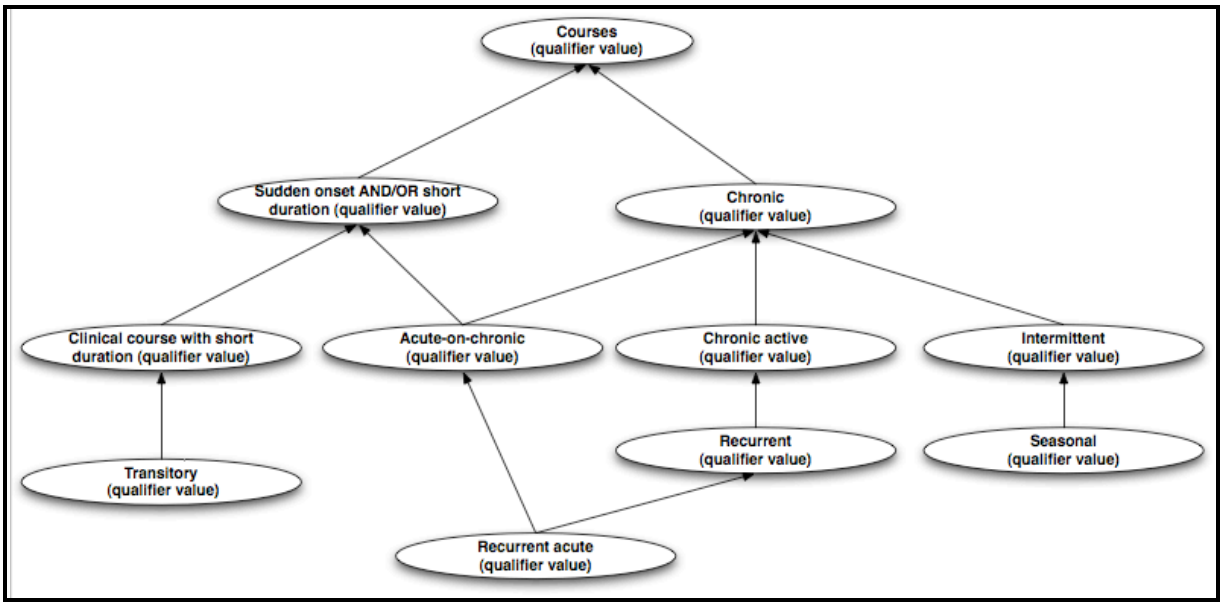

Figure 10: Hierarchy of clinical course qualifiers (original screen shot in on-line appendix)

[28] 405719001|Chronic ulcer (morphologic abnormality)|
[29] 420226006|Pressure ulcer (morphologic abnormality)|
[30] 418172001|Deubitus ulcer (morphological abnormality)|
[31] 62814004|Chronic ulcerative inflammation (morphologic abnormality)|
[32] 405719001|Chronic ulcer (morphologic abnormality)|
[33] 84499006|Chronic inflammation (morphological abnormality)|
[34] 52612000|Lumbar region back structure (body structure)|
[35] 22577007|Posterior abdominal wall structure (body structure)|
[36] 21522001|Abdominal pain (finding)|

6 Discussion

In order to test the use of common qualifiers and the likely accuracy of post-coordination, a variant of a combined lexical and semantic technique first proposed by some of SNOMED CT's founders [6] was implemented and applied to acute and chronic in a module that covers the UMLS CORE Problem List Subset. It is assumed that most developers or clinicians who were not SNOMED specialists would use the generic pattern used to define *Acute disease* and *Chronic disease*, even though there is mention of *acute* and *inflammation* being treated as morphologies in the Users Guide [30].

In the absence of an extensive corpus of post-coordinated expressions, the misclassification rate of existing pre-coordinated expressions as formulated by SNOMED authors was used as a proxy for the likely behaviour of informed clinicians performing post-coordination. Because this is only a proxy, the numerical results must be viewed with caution. At best they are rough approximations of what might take place in real use. All denominators are simply numbers of concepts in SNOMED that meet the lexical criteria; no account is taken for frequency of use although they all occurred in the module extracted from the CORE Problem List Subset, and so can be assumed to be reasonably common. However, the fact that following the pattern given for *Acute disease* and *Chronic disease* produced a misclassification rate on the order of a quarter of concepts tested suggests that the error rate for users would be significant and almost certainly unacceptable.

Of the misclassifications found, only a gratifying small number were due to simple omissions or errors in SNOMED itself. The remainder hinge on whether "acute" and "chronic" should refer to the *clinical course* or *morphology* or both. Representation as *morphology* is clearly appropriate and intuitive to clinical users in a few cases – e.g. leukemias. In most others, it seems arbitrary. The authors could find no list comparable to that in Figure 7b in any of SNOMED's documentation, without which anyone trying to form a post-coordinated expressions would have little chance of avoiding errors. Paradoxically, determining which pattern to use is made even more confusing for users by the fact that some disorders are represented as having both an acute morphology and an acute clinical course, either in their definitions or by inheritance – e.g. *Acute myocardial infarction,*[37] *Acute bronchopneumonia,*[38] and *Acute interstitial pneumonia.*[39] (There is no code for "acute pneumonia" *per se*.) Can users be expected to understand or remember why the *Acute pneumonias* fall under *Acute disease* and *Acute peptic ulcers* do not?

A general solution requires applying general rules and guidelines analogous to the experimental axioms in this paper, plus adding the means to enforce them. To do so would probably require some re-organisation of the *qualifier* and *morphologic abnormality* hierarchies to deal with the issues raised in Section 5.2 and Figure 10 above. The alternative is a case-by-case analysis, which is likely to prove labor intensive and still leave users with distinctions that they find unintuitive and lead to errors. At a minimum, the issue of *morphology* vs *clinical course* must be considered as a "pitfall of post-coordination" and much more clearly documented. How many other such pitfalls exist? The answer remains a topic for future investigations.

The study also demonstrates the usefulness of module extraction and combined lexical and semantic scripting. Module extraction reduced the bulk of SNOMED CT, in this case from 400,000 to 35,000 concepts, reducing the time for classification correspondingly. Without modularisation, the analyses here would have taken many weeks and perhaps been declared

[37] 57054005|Acute myocardial infarction (disorder)|
[38] 123587001|Acute bronchopneumonia (disorder)|
[39] 236302005|Acute interstitialpneumonia (disorder)|

impractical. Focusing on the module derived from the UMLS CORE Problem List Subset, which is claimed to cover ninety-five per cent of SNOMED's use by volume [16] but comprises only ten per cent of its bulk, identified over forty percent of the issues. It also provided experience that could be used to address the remainder efficiently. Scripting allowed the same tests and changes to be applied reproducibly either to modules or the whole.

The total effort required was less than two person-weeks including all experimentation and development – much less than anticipated. The checking, analysis, and discussions concerning the issues raised is taking much longer. However, this experience demonstrates that the technical analyses for quality assurance of SNOMED hierarchies and definitions need not consume excessive resources.

7 Conclusion

Overall, the results suggest that anyone attempting to use post-coordination for common qualifiers such as acute and chronic should be aware that there are serious pitfalls. They also suggest the need to reconcile SNOMED CT's historical legacy in pathology with common usage in patient care. If the goal is to achieve systems that can be used and implemented reliably by those who are not SNOMED specialists, should pathology or the consequences for use in patient care be the deciding criterion?

Acknowledgements

This work supported in part by Siemens Health Solutions. The team working on the Manchester-Siemens collaboration is gratefully acknowledged, as is the input of the WHO ICD-11 revision team. The authors are also grateful to the IHTSDO for making available the stated form of SNOMED CT and the Perl script to convert it to OWL. Particular thanks to Robert Stevens for stimulating this study.

References

[1] International Health Terminology Standards Development Organisation (IHTSDO), Home page, http://www.ihtsdo.org/ (accessed 13 February 2011).

[2] M. Mosquera, Health IT standards panel advances SNOMED use, Government Health IT (2009) http://govhealthit.com/news/health-it-standards-panel-advances-snomed-use (accessed 21 Sept 2011).

[3] US Food and Drug Administration, Indexing structured product labeling for human prescription drug and biological products: request for comments, Federal Register 75 (2010) 33312-33313 http://www.fda.gov/downloads/ForIndustry/DataStandards/StructuredProductLabeling/UCM217297.pdf&rct=j&q=snomed%20fda&ei=u8l5TtXiIcjW0QW5-d3MAQ&usg=AFQjCNF-ArdFoPllntDCeWf0CYBE1Tk2Ig&sig2=KhZbQux6Ty5WhV80P-48Gw&cad=rja.

[4] US Food and Drug Administration, FDA advances federal E-Health effort, Press release, P06-61, (2006) http://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/2006/ucm108642.htm (accessed Sep 2011).

[5] NHS Connecting for Health, SNOMED CT, (2011) http://www.connectingforhealth.nhs.uk/systemsandservices/data/uktc/snomed (accessed Sept 2011).

[6] IHTSDO, Members of IHTSDO, http://www.ihtsdo.org/members (accessed Sept 2011).

[7] J. Rogers and A. L. Rector. Terminological Systems: Bridging the generation gap, in: D. R. Masys (Ed.), Proc. AMIA Fall Symposium (1997) pp. 610-614.

[8] A. L. Rector and J. E. Rogers. 2006 Ontological and practical issues in using a description logic to represent medical concept systems: Experience from GALEN, in: P. Barahona, F. Bry, E. Franconi, N. Henze, and U. Sattler (Eds.), Reasoning Web, Springer-Verlag, Heidelberg, 2006, pp.197-231

[9] Y. Chen, H. H. Gu, Y. Perl, and J. Geller, Structural group-based auditing of missing hierarchical relationships in UMLS, J. Biomed. Inform. 42 (2009) 452-467.

[10] D. Wei and O. Bodenreider. Using the abstraction network in complement to description logics for quality assurance in biomedical terminologies - A case study in SNOMED CT, in: C. Safran, S. Reti, and H. F. Marin (Eds.), Proc. MEDINFO 2010 (2010) pp. 1070-1074.

[11] S. Schulz, B. Suntisrivaraporn, F. Baader, and M. Boeker, SNOMED reaching its adolescence: Ontologists' and logicians' health check, Int. J. Med. Inform. 78 (2009) S86-S94.

[12] W. Ceusters, B. Smith, A. Kumar, and C. Dhaen. Ontology-based error detection in SNOMED-CT, in: M. Fieschi, E. Coiera, and Y.-C. J. Li (Eds.), Proc. MEDINFO-2004 (2004) pp. 482-486.

[13] A. Rector, S. Brandt, and T. Schneider, Getting the foot out of the pelvis: Modelling problems affecting use of SNOMED CT hierarchies in practical applications, J. Am. Med. Inform. Assoc. 18 (2011) 432-440.

[14] K. E. Campbell, M. S. Tuttle, and K. A. Spackman. A "lexically-suggested logical closure" metric for medical terminology maturity, in: C. Chute (Ed.), Proc. AMIA Fall Symposium (AMIA 1988) (1998) pp. 785-789.

[15] K. A. Spackman, K. E. Campbell, and R. A. Côté. SNOMED-RT: A reference Terminology for Health Care, in: D. R. Masys (Ed.), Proc. Annual Symposium of the American Medical Informatics Association (AMIA 97) (1997) pp. 640-644.

[16] K. W. Fung, C. McDonald, and S. Srinivasan, The UMLS-CORE project: a study of the problem list terminologies used in large healthcare institutions, J. Am. Med. Inform. Assoc. 17 (2010) 675-680.

[17] IHTSDO, SNOMED CT User Guide (Jan 2011 International Release), IHTSDO (2011) Included in standard release (accessed June 2011).

[18] F. Baader, Calvanese, D., McGuinness, D. L., Nardi, D., and Patel-Schneider, P. F. The description logic handbook, edition 2. Cambridge University Press, (2007).

[19] W3C OWL Working Group, OWL 2 Web Ontology Language: Document Overview, W3C (2009) http://www.w3.org/2007/OWL/wiki/Document_Overview (accessed 2 Oct 2010).

[20] F. Baader, S. Brandt, and C. Lutz. Pushing the EL Envelope, in: (Eds.), Proc. Nineteenth International Joint Conference on Artificial Intelligence IJCAI-05 (2005) pp. 364–369.

[21] Motic, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokque, A., and Lutz, C. (Eds), OWL 2 Web Ontology Language Profiles, W3C (27 Oct 09) http://www.w3.org/TR/owl2-profiles/ (accessed 5 Mar 2011).

[22] D. Tsarkov and I. Horrocks, FaCT++ description logic reasoner: System description, Automated Reasoning (2006) 292-297.

[23] E. Sirin, B. Parsia, B. Cuenco Grau, A. Kalyanpur, and Y. Katz, Pellet: A practical OWL-DL Reasoner, J. Web Semantics 5 (2007) 51-53.

[24] M. J. Lawley. Exploiting fast classification of SNOMED CT for query and integration of health data, in: R. Cornet and K. A. Spackman (Eds.), Proc. KR-MED 2008 (2008) pp. 8-14.

[25] M. Horridge, S. Bechhofer, and O. Noppens. Igniting the OWL 1.1 Touch Paper: The OWL API, in: C. Golbreich, A. Kalyanpur, and B. Parsia (Eds.), Proc. OWL Experiences and Directions (OWLED 2007) (2007) pp. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.97.4920&rep=rep1&type=pdf.

[26] M. Horridge and P. F. Patel-Schneider. Manchester OWL Syntax for OWL 1.1, in: C. Dolbear, A. Ruttenberg, and U. Sattler (Eds.), Proc. OWL Experiences and Directions (OWLED-2008) (2008) pp. http://www.webont.org/owled/2008dc/papers/owled2008dc_paper_11.pdf.

[27] R. Côté and Rothwell, D. SNOMED-International. College of American Pathologists, Chicago (1993).

[28] K. A. Spackman. Role grouping as an extension to the description logic of Ontylog motivated by concept modelling in SNOMED, in: I. Kohane (Ed.), Proc. AMIA Fall Symposium (AMIA-2002) (2002) pp. 712-716.

[29] R. Cornet and S. Schulz. Relationship Groups in SNOMED CT, in: K. P. Adlassnig, B. Biobel, I. Masic, and J. Mantas (Eds.), Proc. MIE-2009: Medical Informatics in a United and Healthy Europe (2009) pp. 223–227.

[30] IHTSDO, SNOMED Clinical Terms User Guide, IHTSDO (Bundled with SNOMED Release) (2008) Bundled with SNOMED Release (accessed 31 Jan 2011).

[31] B. C. Grau, I. Horrocks, Y. Kazakov, and U. Sattler, Modular reuse of ontologies: Theory and practice, J. Artificial Intelligence Research 31 (2008) 273-318.

[32] L. Iannone, A. Rector, and R. Stevens. Embedding knowledge patterns into OWL, in: L. Aroyo, P. Traverso, F. Ciravegna, P. Cimiano, T. Heath, E. Hyvonen, R. Mizoguchi, E. Oren, M. Sabou, and E. P. B. Bontas (Eds.), Proc. European Semantic Web Conference (ESWC 2009) (2009) pp. 218-232.

[33] K. Dentler, R. Cornet, A. ten Tije, and N. de Keizer, Comparison of reasoners for large ontologies in the OWL 2 EL Profile, Semantic Web 2 (2011) 71-87.

[34] N. Drummond and Manchester Description Logic Group, OWL Patch, (2010) http://owl.cs.manchester.ac.uk/patch/ (accessed Oct 2010).

[35] IHTSDO, SNOMED CT Style Guide: Clinical Findings, IHTSDO (2008) http://www.ihtsdo.org/fileadmin/user_upload/Docs_01/Copenhagen_Apr_2008/SNOMED_CT_Style_Guides/IHTSDO_Modeling_StyleGuide-ClinicalFindings-20080415_v1-07.pdf (accessed 21 Jan 2011).