

Axiomatizing SNOMED CT Disorders: Should there be Room for Interpretation?

Werner CEUSTERS^{a,1} and Anuwat PENGPUT^a

^a*Department of Biomedical Informatics, Jacobs School of Medicine and Biomedical Sciences, University at Buffalo, Buffalo NY, USA.*

ORCID ID: Werner Ceusters <https://orcid.org/0000-0002-2676-8689>,
Anuwat Pengput <https://orcid.org/0000-0002-0273-1531>

Abstract. SNOMED CT is a large *concept*-based terminology designed according to epistemic, semantic and pragmatic principles relevant to clinicians. Its goal is structured clinical reporting in electronic healthcare records (EHRs). The Basic Formal Ontology (BFO) is an ontology designed on the basis of *types* claimed to exist in reality based on a domain-independent ontological theory. Its goal is faithful representation of reality within that theory. The Ontology for General Medical Science (OGMS) extends the BFO by providing definitions for types relevant within the clinical domain. Combining SNOMED CT with the ontological rigor of BFO and OGMS might improve clinical reporting by, *f.i.*, preventing data entry mistakes and inconsistencies, and make EHRs more comparable. To that end, we are developing a *logical* framework capable of exploiting what SNOMED CT offers *terminologically* and realism-based ontologies such as the BFO and the OGMS *ontologically* by means of bridging axioms compatible with the BFO, and expressed in the same CLIF-dialect as used in its axiomatization in first order logic. In this paper, we report on our attempts to detect in the combinations of binary relations that are used in the definition of SNOMED CT's definitions of disorder concepts patterns which might at least partially automate the construction of such axioms. Our findings suggest that this partial automation is indeed possible, but to a smaller extent than we had hoped for. We compare our approach with a recent proposal that seeks to bring SNOMED CT and BFO closer together by reinterpreting SNOMED CT disorders as clinical occurrents. The proposal has its merit in providing a realist underpinning for that part of SNOMED CT's concept model in terms of the BFO, but is not discriminatory enough for an automatic translation into OGMS. Key problem is the lack of face validity of SNOMED CT disorder terms as compared to the formal definitions they are given and this in absence of textual definitions.

Keywords. Applications and Methods, SNOMED CT, Basic Formal Ontology, CLIF

1. Introduction

SNOMED CT is a large clinical terminology designed to enable consistent representation of clinical content in electronic health records (EHR) and is claimed to be thereto used in over 80 countries world-wide [1]. Terms in SNOMED CT are given meaning through *concepts* as intermediary devices organized in a taxonomy inferred by a classifier on the basis of concept definitions expressed in the description logic EL⁺⁺ [2]. Several studies

¹ Corresponding Author: Werner Ceusters, Department of Biomedical Informatics, University at Buffalo, 77 Goodell Street, Buffalo NY, 14203, USA. Email: wceusters@gmail.com

over the past 20 years indicate that SNOMED CT's concept coverage is generally found to be excellent for its design purpose and that it performs very well in semantic searches over patient records for cohort formation in clinical trials or for assessments of quality of care within certain populations [3]. Of course, in light of its size – the SNOMED CT International version of November 2022 which we used for the analysis presented in this paper (from here on abbreviated as 'sctiv2211') counts 358,723 active concepts, 576,439 active '*isa*' relationships and 639,483 associative relationships – it still contains mistakes and inadequate classifications primarily due to underspecified concept definitions. Although these problems are sometimes hard to find [4, 5], they can be fixed easily. Yet, to realize its full potential, at least two major barriers should be overcome. The first one is the hitherto inadequate integration of SNOMED CT and various other terminologies and classification systems into EHRs; this is known as the 'terminology binding problem' [6]. It causes various misalignments between patients' physical realities, clinicians' mental models, and structured documentation in EHRs, resulting in several types of misrepresentation leading to human error [7]. The second barrier is SNOMED CT's ontological commitment. It has variably been described as lacking, inconsistent and ambiguous, thus leaving open the question what SNOMED CT terms actually denote [8-11]. Furthermore, none of the upper and higher middle level concepts in terms of which SNOMED CT's concept model is constructed or which are the most general subsumers of semantic categories such as *disorder*, *clinical finding*, *event*, *procedure*, etc, are formally defined. This hampers clinical decision support, re-use of data for research purposes, and automated quality control for data input in EHRs.

The Basic Formal Ontology (BFO), in contrast, and a few of the many ontologies derived therefrom such as the Ontology for General Medical Science (OGMS) [12], have a very precisely defined ontological commitment: terms therein are not organized on the basis of meanings but on the basis of a domain-independent ontological theory [13]. BFO's and OGMS's goal is to represent reality faithfully in line with that theory and independent of any potential use. But both of them being upper level ontologies – BFO independent of any domain, OGMS defining only the most general subsumers of types relevant to biomedicine – they are tiny. Their potential can only be realized when used to aid the development of new domain ontologies, or when linked or integrated in one or other form with existing ontologies. For clinical medicine, SNOMED CT is an obvious choice.

A few attempts have been made to link BFO and OGMS to SNOMED CT. In one study it was concluded to be impossible because of the mutual incompatibility of the different underlying categorization principles [14], while in another one a concrete mapping was proposed [15]. The latter proposal is on closer inspection seriously flawed, in part due to misunderstandings about the underlying models on either side and in part due to the limitations of the description logic used for the mapping. Recently, Schulz *et al.* saw room for a partial integration of BFO and SNOMED CT by reinterpreting SNOMED CT's *clinical finding* concepts, which include *diseases*, as denoting something that under BFO would be an *occurrent* [16]; the SNOMED CT term 'Cholangiocarcinoma of the biliary tract' would then not mean what one would assume it means, but rather '*having* a cholangiocarcinoma of the biliary tract'!

In this paper, we elaborate further on our approach to let SNOMED CT's view and BFO's view happily co-exist by means of bridging axioms in first-order logic (FOL) that combine SNOMED CT's view with the perspective adhered to by the BFO and the OGMS, and this without violating the principles adhered to by either side [17]. Here we explore specifically the extent to which SNOMED CT's concept model for disorder

definitions can be used to automate – at least partially – the creation of such axioms and the potential impact of Schulz *et.al.*'s proposal on this approach.

2. SNOMED CT

SNOMED CT is built out of four core components: terms, concepts, descriptions and relationships. Every *term* in SNOMED CT, e.g. 'Cholangiocarcinoma', is linked to at least one *concept* by means of at least one *description*. According to the custodian of SNOMED CT, SNOMED International, 'Every concept represents a unique clinical meaning' [1]. Terms can be linked to more than one concept, what accounts for homonymy. Multiple terms can be linked through descriptions to some unique concept, thus allowing for synonymy. All but one concepts are linked to other concepts by means of concepts called *attributes* which are used to express *relationships*. These relationships come in two flavors: the ones created by means of the attribute *isa* express subsumption, while many other attributes express various sorts of associations. Although there are 1191 associative attributes listed in sctiv2211, only 99 (8%) are used in concept definitions, with a total of 36,579 (10%) unique concepts as range. As an example, Figure 1, retrieved from the browser through which sctiv2211 was fully accessible [18], shows for the concept with unique identifier '312104005' that it has two taxonomic parents, four children, and two associative relations, one formed by means of the attribute 'Finding site' and one by means of the attribute 'Associated morphology'.

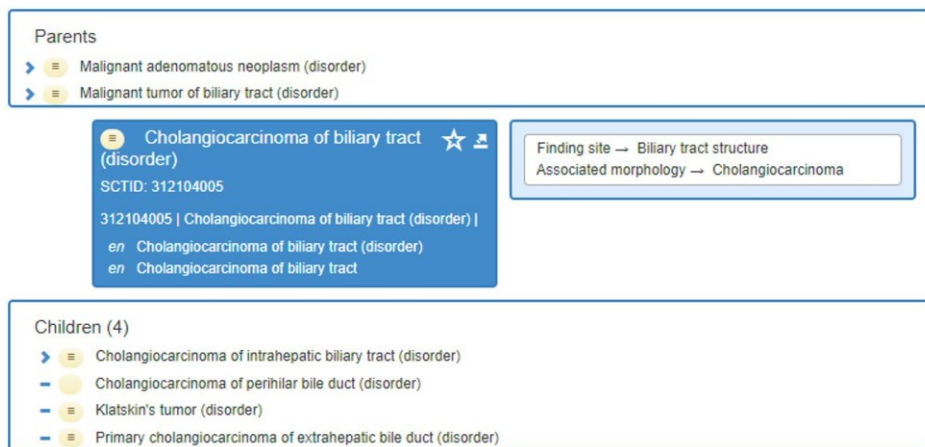


Figure 1. Terms and relationships linked to SNOMED CT concept 'sctid:312104005'.

2.1. Meanings through Fully Specified Names

Each concept is linked to precisely one term designated as the *fully specified name* (FSN) and which is claimed to 'represent a unique, unambiguous description of the concept's meaning' [1]. All FSNs come with a *semantic tag* at the end. For the concept in Figure 1, that is 'disorder'. Semantic tags – there are 58 different ones in sctiv2211 – are in the first place intended to clarify the *intended meaning* of concepts that have terms that stand in a homonymous relation to each other as f.i. in 'Temperature (observable entity)', 'Temperature (attribute)', 'Temperature (property) (qualifier value)', but also closely

related: ‘*Low temperature (physical force)*’ and, surprisingly, ‘*Food temperature (substance)*’. Some of the tags indicate what major subsumption hierarchy the concept described by it belongs to, while others play a role in SNOMED CT’s concept model which specifies the domain-range restrictions that must be applied to attributes. For almost every semantic tag, there exists precisely one concept that bears that semantic tag and that subsumes all other concepts with that tag. When that is not the case, it is most likely either a mistake (f.i. 7 of the 709 concepts with the tag ‘cell’ are for no good reason not subsumed by the concept ‘*Entire cell (cell)*’), or, at least for the time being, to be considered a design choice: for the 7819 concepts with the tag ‘*clinical drug*’, there is not one most general subsumer with that tag.

2.2. Meanings through formal language

While the FSN is intended to provide a human-readable description of the meaning of the concept to which it belongs, the relationships the concept enjoys are intended ‘*to logically define the meaning of a concept in a way that can be processed by a computer*’ [1]. SNOMED CT authors craft the definitions as dictated by the concept model [19]. Some of these definitions are available in an OWL 2 dialect which in `sciv2211` is restricted to 10 constructs with `EquivalentClasses`, `ObjectIntersectionOf`, `ObjectSomeValuesFrom`, and `SubClassOf` being the only ones used for disorder definitions. 36.5% of SNOMED CT’s concepts are defined by means of the `EquivalentClasses` construct.

Unfortunately, none of the high-level concepts that are the subsumers of the main hierarchies, or are used to specify domain-range constraints for attributes, are formally defined. Neither are the attributes themselves; for some of them, there is a description attached to the concept that provides some explanation, but the majority require consulting the documentation [19]. Only 4 attributes enjoy the `TransitiveObjectProperty`. Three of them express parthood but are not used in any class expression. Parthood is in SNOMED CT thus far handled by structure-entirety-part (SEP) triplets [20] such that for each entity X for which parthood representation is required, three concepts are created and related so that ‘structure of X’ subsumes both ‘entire X’ and ‘part of X’.

As an example of a disorder definition, expression (1) gives the OWL-expression in SNOMED CT’s own functional-style syntax for the concept of Figure 1 [18]:

```
EquivalentClasses(
:312104005 |Cholangiocarcinoma of biliary tract (disorder)|
ObjectIntersectionOf(:64572001 |Disease (disorder)|
ObjectSomeValuesFrom(:609096000 |Role group (attribute)|
ObjectIntersectionOf(
ObjectSomeValuesFrom(
:116676008 |Associated morphology (attribute)|
:70179006 |Cholangiocarcinoma (morphologic abnormality)|))
ObjectSomeValuesFrom(:363698007 |Finding site (attribute)|
:34707002 |Biliary tract structure (body structure)|)))))) (1)
```

Not provided in OWL, but derivable from the *isa*-relationships, are the following two expressions for that concept:

```
SubClassOf( :312104005 |Cholangiocarcinoma of biliary tract (disorder)|
:443961001 |Malignant adenomatous neoplasm (disorder)|) (2)
```

SubClassOf(:312104005 |Cholangiocarcinoma of biliary tract (disorder)| (3)
:363415003 |Malignant tumor of biliary tract (disorder)|)

Important here, as exemplified in expression (1), is the use of the ‘*Role group (attribute)*’ concept. It was originally introduced to prevent SNOMED CT’s proprietary OntyLog classifier from misclassifying concepts with definitions in which an attribute appeared more than once while linked to different target concepts [21]. Its use was later expanded to give certain concepts an inclusive reading so that, for example, a ‘*fracture of radius and ulna*’ would be subsumed by both ‘*fracture of radius*’ and ‘*fracture of ulna*’ [16]. Role groups are used in concepts with one out of only 12 of the 58 semantic tags, but nevertheless account for almost 62% of all concepts. All concepts with the tags ‘*clinical drug*’ and ‘*medicinal product form*’ have at least one role group and are all what SNOMED CT calls ‘fully defined’, i.e. are defined with necessary and sufficient conditions. That is, oddly, not the case for 8 of the 8413 ‘*medicinal product*’ concepts, of which 3 are defined with only necessary conditions. Other heavily role-grouped concepts are those with the tag ‘*disorder*’ (2 of which enjoy 13 role-groups), ‘*specimen*’, ‘*situation*’ and ‘*procedure*’.

2.3. Disorder concepts in SNOMED CT

SNOMED CT is designed on the basis of clinical practice-oriented considerations as for example pragmatic ones: the taxonomy is so structured that clinicians should easily find their way in it. It has also a strong epistemic basis as witnessed by the main taxonomic hierarchies *clinical finding* and *observable entity* which together account for nearly 30% of the content. What it is for SNOMED CT to be a *clinical finding*, is to be ‘*the active acquisition of subjective or objective information from a primary source. This includes information acquired from human observers, through recording of data via the use of scientific instruments, or indirectly from samples taken from the source, and evaluated separately*’ [19, p159]. To be an *observable entity*, is to be ‘*information about a quality/property to be observed and how it will be observed*’ [19, p14] and/or ‘*the name of something that can be observed and represents a question or assessment which can produce an answer or result*’ [19, p159]. For example: ‘*color of nail is an observable entity*’ and ‘*gray nails is a finding*’ [19, p244, bold emphasis added].

In SNOMED CT, all disorders are subsumed by *clinical finding*. This is defensible in terms of the purpose for which SNOMED CT is designed, i.e. clinical documentation: clinicians can only assert patients to be diseased if they found that to be the case. But in light of the description of clinical finding provided in the documentation, this is rather awkward as it would entail that in the eyes of SNOMED CT disorders are active acquisitions of information what is contradicted by the properties ascribed to disorders in the very same document: ‘*Always and necessarily abnormal; necessarily have an underlying pathological process; have temporal persistence (may be under treatment, in remission, or inactive, even though they are still present); may be present as a propensity for certain abnormal states to occur, even when treatment mitigates or resolves those abnormal states*’ [19, p160].

It is the latter view, i.e. that disorders are entities on the side of the patient and not processes of information gathering, that is coherent with the list of attributes that are used in SNOMED CT to define disorders (Table 1). But therein also is it easy to identify oddities brought about by epistemic considerations. *Cholangiocarcinoma of biliary tract*, for example, is according to expression (1), and contrary to what one might expect on

the face value meaning of the term, not subsumed by *cholangiocarcinoma*, the latter being a *morphologic abnormality* and the former a *disorder*. Whereas ‘*The concepts in the morphologic abnormality hierarchy represent abnormal body structures*’ [19, p153], subtypes of *neoplasm (morphologic abnormality)* ‘*represent histological cell types that are recognized internationally by pathologists, classified by WHO, and aligned with the ICD-O classification*’ [19, p156]. So here again SNOMED CT posits distinctions on entities on the side of the patient not on the basis of what these entities intrinsically are, but on how or by whom they are described in other systems.

Table 1. Associative attributes in SNOMED CT concepts subsumed by *disease (disorder)*.

Attribute	Occurrence	Explanation
Associated morphology	67.22%	specifies the morphologic changes seen at the tissue or cellular level that are characteristic features of a disease.
Associated with	1.42%	represents a clinically relevant association between concepts without either asserting or excluding a causal or sequential relationship between the two.
Causative agent	18.69%	Identifies the direct causative agent of a disease such as an organism, substance or physical force. (Note: This attribute is not used for vectors, such as mosquitos transmitting malaria).
Clinical course	6.27%	represents both the onset and course of a disease.
Due to	20.18%	relates a clinical finding directly to a cause such as another clinical finding or a procedure.
Finding informer	<0.1%	specifies the person (by role) or other entity (e.g. a monitoring device) from which the clinical finding information was obtained. This attribute is frequently used in conjunction with finding method.
Finding method	<0.1%	specifies the means by which a clinical finding was determined. This attribute is frequently used in conjunction with finding informer.
Finding site	81.12%	specifies the body site affected by a condition.
Has interpretation	6.83%	when grouped with the attribute interprets, designates the judgment aspect being evaluated or interpreted for a concept. (e.g. presence, absence etc.)
Has realization	<0.01%	specifies the realization of a function
Interprets	9.26%	refers to the entity being evaluated or interpreted, when an evaluation, interpretation or judgment is intrinsic to the meaning of a concept.
Occurrence	17.53%	refers to a specific period of life during which a condition first presents.
Pathological process	24.53%	provides information about the underlying pathological process of a disorder, i.e. it describes the process that results in the structural or morphologic change.
Severity	<0.01%	used to sub-class a clinical finding concept according to its relative severity.
Temporally related to	0.06%	a period of time occurring before, during and or after a clinical entity. Subsumes <i>after</i> (3.29%), <i>before</i> (< 0.01%) and <i>during</i> (0.59%).

3. The Basic Formal Ontology

The BFO differs in important aspects from SNOMED CT as it follows an ontological theory which rests on the distinctions between (1) *types* and *particulars*, (2) *dependent* and *independent* entities, (3) *continuants* and *occurrents*, and (4) *representations* and *what representations are about* [13]. Although there are a few types in BFO which carry the same label as some concepts in SNOMED CT, e.g. *disposition* and *function*, there is absolutely no correspondence between them. To avoid confusion, we will use from here on the prefix ‘*bfo:*’ to indicate a reference to a BFO type, and ‘*sct:*’ when referencing either an individual concept such as *sct:disease (disorder)*, or any or all concepts that

carry some semantic tag, such as *sct:disorder* and *sct:disorders* respectively. Then we can state f.i. that all *sct:disorders* are subsumed by *sct:clinical finding (finding)*.

Extremely relevant for the case we try to make here is that important differences between BFO and SNOMED CT become blurred when BFO is looked at through its OWL representation. In OWL, BFO's *types* become rendered as *classes*. However, none of these classes should be understood as being formed on the basis of properties that individuals of these classes have 'in common'. To render its taxonomy, BFO-OWL uses axioms of the form *SubClassOf(:A :B)*, but, by doing so, two important aspects of the ontology are lost. The first one is that A is not just *more specific* than B – i.e. the rather broad OWL meaning of *SubClassOf*, typical for the concept-based approach [22] – but only in terms of the ontological *instantiation-at-a-time*-relation between *particulars* and *types*, not to be confused with OWL's set-theoretic meaning of *individuals* being instances of class expressions. To clarify this point, imagine that some realism-based ontology (say RBO) which adheres to BFO's principles includes a reference to the type *rbo:fracture of radius and ulna (disorder)* (if indeed a case can be made for the existence of a type corresponding to *sct:fracture of radius and ulna (disorder)* vs. such expression simply stating in a terminologically concise manner the existence of two distinct entities). Based on BFO's principles, RBO cannot accept that an instance of *rbo:fracture of radius and ulna (disorder)* would also instantiate *rbo:fracture of radius* and *rbo:fracture of ulna*: two things cannot be a kind of one thing. SNOMED CT contained in the past many assertions of this sort, e.g. '*both testes isa testis*', thus exhibiting *isa-overload* [23]. The obviously odd ones have since then be removed, except for combinations occurring in syndromes, or to make retrieval more easy as in the double fracture case [16].

The second aspect lost is that BFO's instantiation, as well as many other relations, is time-indexed: for the BFO, the assertion *SubClassOf(:A :B)* implies that it must be the case that *at every time* a particular instantiates type A, it instantiates also type B *at that time*. Loss of time-indexing is the result of OWL-DL's inability to express relationships with more than two arguments. Time indexing, however, applies to all relationships in which a continuant is involved, and is, since BFO version *BFO2020* is accepted as an ISO standard, specified explicitly in its axiomatization in first-order logic [24], as exemplified in expression (4) in Common Logic Interchange Format (CLIF).

(forall (a b) (iff (inheres-in a b) (4)
 (and (specifically-depends-on a b)
 (exists (t) (and (instance-of a specifically-dependent-continuant t)
 (instance-of b independent-continuant t)
 (not (instance-of b spatial-region t))))))))

We therefore argue that to use BFO optimally, one should not work with the OWL-version, but with the FOL-version. After all, it has already been shown that by integrating BFO's FOL axioms in the Ontology of Biomedical Investigations, one of the very few OBO Foundry ontologies that adhere to BFO, inconsistencies could be detected [25].

4. The Ontology for General Medical Science

The OGMS was created to avoid the confusions that are exhibited in other biomedical ontologies and terminologies, including SNOMED CT, in which diseases are conceptualized as findings or forms of evidence, and in which diseases are not

distinguished from disorders although the words ‘disease’ and ‘disorder’ have distinct collocational patterns in many languages [12]. In OGMS, diseases are defined as dispositions rooted in physical disorders and realized in pathological processes. This approach allows for a cleaner formal representation of pre-clinical manifestations of disease, and of combinations of disease and predispositions to disease which can exist within a single patient. As will be further demonstrated in section 5, there is no correspondence between *ogms:disorder* and *sct:disease (disorder)*. Some *sct:disorders* correspond to *ogms:disorders* while others correspond to *ogms:diseases*, *ogms:disease courses* or *ogms:pathological processes*.

Table 2 contains the most relevant definitions for the topic covered here. OGMS is thus far not axiomatized in first-order logic.

Table 2. Core definitions in the Ontology for General Medical Science

Etiological Process =def.	– A process in an organism that leads to a subsequent disorder.
Disorder =def.	– A causally relatively isolated combination of physical components that is (a) clinically abnormal and (b) maximal, in the sense that it is not a part of some larger such combination.
Pathological Process =def.	– A bodily process that is a manifestation of a disorder.
Disease =def.	– A disposition to undergo pathological processes that exists in an organism because of one or more disorders in that organism.
Disease Course =def.	– The totality of all processes through which a given disease instance is realized.

5. Bringing SNOMED CT and BFO together

Although a direct alignment between SNOMED CT’s *concepts* and BFO’s and OGMS’ *types* is not possible, we believe a lot can be gained by combining the two representations for use in clinical applications. By doing so, such applications could benefit from SNOMED CT’s terminological richness: there are in *sciv2211* close to 1.5 million terms linked to the 358,723 concepts, in contrast to the roughly 80 types of BFO and OGMS together which have no other terms than the type labels. With adequate harmonization between SNOMED CT and BFO/OGMS, applications could use the ontological rigor of the latter to identify and prevent data entry mistakes in clinical records. An application could f.i. generate alerts when a clinician asserts in the problem list of an EHR that a patient has ‘prediabetes’ while it is already asserted that he has ‘diabetes’; that a ‘bilateral cerumen impaction’ transitions into a unilateral one [26, 27]; or that a tumor, an instance of *ogms:disorder*, becomes cancer, an instance of *ogms:disease* [28]. Yet, a combined representation, in whatever form, should also ensure that the differing perspectives of the source representations are respected.

5.1. Reinterpreting *sct:disorders* as *bfo:occurents*

In a recent proposal, Schulz *et.al.* argue – insofar we understand them correctly – that if a clinician puts the SNOMED CT term ‘*Cholangiocarcinoma of biliary tract (disorder)*’ on the problem list of a patient’s EHR, the term should not be understood as denoting the cancer in the patient’s bile duct, but an entity of the type ‘*having a cholangiocarcinoma of the biliary tract*’ [16]. Such entities would be *clinical occurents*. Role groups in *sct:disorder* and *sct:finding* definitions would be interpreted as the *bfo:has-occurent-part* relation and *sct:morphological abnormalities* would be subsumed by *bfo:material entity*. As a result, so they argue, *sct:disorders* can be smoothly integrated with BFO without significant change to either SNOMED CT or BFO.

The proposal rests on the in our opinion correct assumption that when an experienced biomedical ontologist adhering to Ontological Realism and a surgeon, whether or not familiar with SNOMED CT, are discussing a concrete medical case, they are referencing in their dialogue the very same entities on the side of the patient, for instance the carcinoma in the patient's biliary tract. It is also totally irrelevant for the surgeon who takes care of the patient, and ideally must get rid of the carcinoma, whether the carcinoma is an instance of *bfo:material entity*, *bfo:disposition*, or *bfo:process*, or of the three closely related types *ogms:disorder*, *ogms:disease*, or *ogms:disease course*. However, the distinctions are relevant for tools that can prevent the sort of issues in clinical records mentioned above. If Schultz *et.al.*'s proposal would be accepted, *sct:disorders* would become *sct:clinical occurrents* and subtypes of *bfo:occurents*. The definiens of expression (1) would then in Manchester OWL syntax become:

'*sct:Cholangiocarcinoma of biliary tract*' equivalentTo (5)
'*sct:clinical occurrent*' and **bfo:has-occurrent-part** some
(('sct:Finding site' some '*sct:Biliary tract structure*') and
('sct:Associated morphology' some *sct:Cholangiocarcinoma*))

The merit of this proposal is that the SNOMED CT authors would recognize the value of the BFO; why otherwise doing this effort? The practical value might however be limited, as it continues to see BFO through its crippled OWL rendering without doing justice to the underlying theory. As we will show below, many SNOMED CT attributes would also need to be reinterpreted to be applicable to types recognized by the BFO. Furthermore, if indeed *sct:disorders* need to be interpreted as *having a disorder*, SNOMED CT would face a conundrum: all these terms would lose face validity and be inconsistent with the new definition of the concept, or, all *sct:findings* need to be made obsolete and replaced by new concepts according to SNOMED CT's stable meaning policy. Neither is a minor change!

5.2. An approach with bridging axioms

We argue that SNOMED CT and BFO/OGMS should not be combined in one *ontological* framework, but in one *logical* framework capable of exploiting what SNOMED CT offers *terminologically* and realism-based ontologies such as the BFO and the OGMS *ontologically* [17]. This framework does also not require changes in BFO or SNOMED CT. It favors an interpretation of SNOMED CT terms which is as closely as possible faithful to what the FSN suggests – this is after all the only information clinicians can see in contemporary 'picking list' or 'value set'-based implementations of SNOMED CT in EHRs – while at the same time taking into account the elements out of which the formal definition is composed. Within this framework, SNOMED CT concepts are considered to be instances of *information content entity*, thus of *bfo:generically dependent continuant*, and *to be about* the classes which form the concepts' extensions [29]. Patient data expressed in terms of SNOMED CT concepts are considered to be instances of *information quality entities*, thus of *bfo:quality*, and *to be about* existing entities on the side of the patient. Here we describe only how the framework is set up concerning the concepts.

BFO's domain of discourse contains individuals unary related as either *universal* (denoting some type) or *particular*. We include in our framework's domain – not in BFO's domain of course – individuals unary related as *class*, the latter corresponding to

the set of all individuals that satisfy a concept's meaning, i.e. the concept's extension. Using the same CLIF-dialect as used for BFO's axiomatization, we can then list explicitly all classes for the concepts we accept in the framework. For all concepts we do accept, we accept also that any individual in the corresponding class is numerically identical with some particular under BFO's perspective. In analogy with the ternary BFO-relation *instance-of* used to assert at what time a particular belongs to some universal, we use the binary relation *individual-of* to assert to what class some particular belongs.

(class sct-cholangiocarcinoma-of-biliary-tract) (6a)
 (forall (x y) (if (individual-of x y) (and (particular x) (class y)))) (6b)

From axiom (6b) and BFO's existence-instantiation axioms, it follows that every individual that satisfies some accepted concept's meaning exists at some temporal region and whenever it exists, instantiates some universal. A possible translation of SNOMED CT's perspective in the perspective of BFO and OGMS would then for the concept defined in expression (1) be achieved through expression (7):

(forall (x y z) (7)
 (if (and (individual-of x sct-disease)
 (sct-finding-site x y)
 (individual-of y sct-biliary-tract-structure)
 (sct-associated-morphology x z)
 (individual-of z sct-cholangiocarcinoma))
 (exists (p t) (and (continuant-part-of p z t)
 (located-in p y t)
 (instance-of z ogms-disorder t)
 (instance-of y ogms-bodily-component t))))))

Expression (7) takes care of the inclusive reading intended by the role group in (1) by not stating how *x* fits in the OGMS perspective. It would also be impossible to do since *sct.disease (disorder)* is not defined in SNOMED CT and individuals in its extension can be instances of distinct OGMS types. It might thus be that for OGMS the including and included part are instances of distinct OGMS types, thus making it impossible to express how they would be related. For an exclusive reading, thus if the role group were not there, and what might f.i. be defined in national extensions to the international version, the conjunction in the consequent part of (7) could be expanded with the clauses shown in (8).

(iff (instance-of x ogms-disease t) (material-basis-of z x t)) (8a)
 (iff (instance-of x ogms-disease-course t) (8b)
 (exists (d) (and (instance-of d ogms-disease t)
 (realizes x d) (material-basis-of z d t))))

Since the translation of EL++ to FOL is relatively straightforward [25], the antecedent part of the bridging axioms can be generated easily. However, because of the different perspectives taken by SNOMED CT and realism-based ontologies, the consequent part requires much more caution. That is at least the case for (1) the number of variables to be used and the types instantiated by what they denote, (2) the temporal

indexing that needs to be applied including situations in which some predicates require temporal indices distinct from those in other predicates, and (3) the different BFO relations that are required because of the different OGMS types *sct:disorders* can fall under. This is because we want in the current setup of our framework all predicates in the consequent part of the axioms to be BFO-based, and not a mixture of SNOMED CT terms and BFO relations as proposed by Schulz *et.al.*

Here we report our results in attempting to detect, in the combinations of binary relations that are predominantly used in the definition of SNOMED CT disorder concepts (Table 1), patterns which might at least partially automate the construction of such axioms. Since SNOMED CT distributions thus far do not contain a complete rendering in OWL, an analysis to determine usable patterns must include also all the triples provided in SNOMED CT's *relationships table*.

5.3. Bridging patterns for *sct:disorders*

All *sct:disorders* are subsumed by *sct:clinical finding (finding)*, i.e. the most global subsumer of the concepts which are the topic of the reinterpretation proposal of Schulz *et.al.* [16]. As part of our strategy, we computed for each *sct:disorder* (1) the number of occurrences of each of the 18 different attributes used in its definition, (2) the number of attribute occurrences it requires through its subsumers, and (3) the number of role groups. We then tried to look for combinations of attribute occurrences, within and across role groups, that are either indicative for an impossibility to provide an OGMS perspective that corresponds with the face value meaning of the disorder term, or suggestive for a high rate of success. We also investigated the effects of the reinterpretation proposal where possible.

5.3.1. Concepts with no attributes

Of the 80,239 *sct:disorders* in *sctiv2211*, including *sct:disease (disorder)* itself, 7,675 (9.6%) do not have any of the 18 attributes listed in Table 1. 2104 thereof do not have any subsumer with attributes either. Obviously, for the latter, no automatic OGMS determination can be made, thus requiring a manual inspection of the FSN to apply the principle of face-validity. For example, the face-validity principle would suggest, in the absence of any other information and based on the definitions in Table 2, 'tumor of ...' to be interpreted as *ogms:disorder* and 'chronic ...' as either *ogms:disease course* or as *ogms:disease bfo:has-realization ogms:chronic-disease-course*. Also under the reinterpretation proposal would a manual inspection be required. This is because the proposal accepts *states* as *sct:clinical occurrents* while BFO currently does not accept 'state' as a type. 'State' as type seems rather to be perceived as needed when one uses a language not capable of time-indexing. Using BFO's CLIF-dialect, the term '(and (instance-of *x* 37°C-temperature *t*) (instance-of *t* temporal-interval *t*))' expresses clearly that *x* remains constant at 37°C precision within the temporal interval *t*.

5.3.2. Finding site and associated morphology

The finding site attribute is used at least once in 66,546 (82.9%) *sct:disorder* definitions and is range-restricted to anatomical structures. When the *sct:disorder* is a *bfo:material-entity*, the attribute may be translated into *bfo:continuant-part-of* when its range is a *bfo:material-entity* too. If the attribute's range is a *bfo:site* such as an armpit, *bfo:located-in* should be used. Since for BFO every *x* which is *bfo:continuant-part-of* some *y* at some

time t is also *bfo:located-in* y at t , the latter seems to be the safer option. Unfortunately, *sct:finding site* is also used to assert that the stated range does not exist at all, as in *sct:congenital complete absence of left upper limb (disorder)*. This is a consequence of the inability to express negation in SNOMED CT's DL. In such cases, the finding site attribute is typically combined in a role group with an additional '*sct:associated morphology sct:agenesis (morphologic abnormality)*' assertion. For the congenital left upper limb absence example, it is no problem to express in the consequent part of the bridging axiom that for all instances of left upper limb and all human beings that have that disorder, there does not exist a time at which such limb instance is part of such human being instance. But this works only for missing body parts of which a 'complete' human being has only one instance. For *sct:agenesis of nerve (disorder)*, for example, the 'nerve' variable needs to be existentially quantified. Since SNOMED CT does currently not encode explicitly of which anatomical structure concepts 'complete' human beings can have only one individual, we are again left with manual inspection only. Caution is still required when location or parthood are applicable and this because of the vagueness and impreciseness of *sct:finding site*. Hence the introduction of p in the consequence part of expression (7): it might very well be that a patient's cholangiocarcinoma breaks through the boundaries of what qualifies as the biliary tract so that simply asserting continuant parthood between z and y would be wrong in such case. But when it happens, there would still be a part p of the carcinoma which is located in the biliary tract at time t while other parts are outside the biliary tract.

When on face-value the *sct:disorder* is a *bfo:process*, then *sct:finding site* might correspond to either *bfo:occurs-in* or *bfo:has-participant*. Both relations work for *bfo:sites* and *bfo:material-entities* as the *sct:finding sites*. When the former is used, one would express only occurrent parthood of the spatiotemporal region wherein the process evolves and the spatiotemporal region in which the *sct:finding site* exists. With the latter, one would assert that the process cannot exist without the existence of the *sct:finding site*. Which of the two applies in any specific case, if not both, or perhaps none at all, requires not only medical expertise but also clear definitions. Therein falls SNOMED CT short as exemplified by *sct:atrial fibrillation (disorder)* which is defined as being an *sct:fibrillation (disorder)* with two finding sites: *sct:atrial structure* and *sct:cardiac conducting system structure*. Yet, these two sites are differently related to the fibrillation, a difference SNOMED CT fails to express by conflating distinct OGMS types, but which is appropriately covered in expression (9): it is the cardiac atrium that participates in a fibrillation which realizes the disease which has as physical basis a disorder in the cardiac conduction system.

(forall (x y z) (9)
 (if (and (individual-of x sct-fibrillation)
 (sct-finding-site x y)
 (individual-of y sct-atrial-structure)
 (sct-finding-site x z)
 (individual-of z sct-cardiac-conducting-system-structure))
 (exists (do di t1 t2) (and (instance-of do ogms-disorder t1)
 (continuant-part-of do z t1)
 (instance-of di ogms-disease t1)
 (material-basis-of do di t1)
 (instance-of x ogms-pathological-process t2)
 (realizes x di) (has-participant x y t2))))))

5.3.3. Pathological process

This attribute is used in the definition of 17,026 (21%) *sct:disorders*, accounts for about 10% of all relationships with as domain *sct:disorder*, and has as range only 19 distinct concepts all of which being subsumed by *sct:qualifier value*, i.e. the top of a hierarchy with 11,727 concepts. Qualifier values fall apart in two groups. The one relevant here consists of only 875 (7.5%) concepts and is described as to ‘include a wide range of concepts that provide attribute values used in the definitions of other concepts’ [19, p470], a rather odd characterization in light of the 35,704 other concepts that are used as range in concept definitions and belong to other hierarchies. The *sct:qualifier values* used as value for *sct:pathological process* are all related to hypersensitivity including allergy, with five exceptions: maldevelopment, malignant proliferation of a primary neoplasm, inflammation, infection and parasitic growth. For the hypersensitivity related ones, it turns out that the *sct:disorder* being defined can be any of the five OGMS types listed in Table 2, thus requiring manual inspection in every single case. Although all five exceptions correspond to *bfo:process*, the former four to *ogms:pathological process* and the latter one to *ogms:etiological process*, solid patterns could thus far not be found either. There is to a certain extent the combination with *sct:agenesis (morphologic abnormality)* discussed in section 5.3.2. Because the OGMS is currently not yet axiomatized, it seems relatively safe to equate *sct:disorders* with *ogms:disease course* when the following conditions are met: (1) there is only one role group, (2) formed by combining ‘*sct:pathological process (attribute) sct:infectious process (qualifier value)*’ with *sct:causative agent (attribute)* with as range any virus or bacterium concept, and (3) in absence of any *sct:finding site (attribute)* in the definition. This is motivated as follows. An infectious disease in a patient may of course involve distinct body parts belonging to distinct bodily systems. A patient with tuberculosis might have at the same or at different times lung tuberculosis and bone tuberculosis. Both that patient’s lung and bone tuberculosis are *ogms:pathological processes*, i.e. manifestations of the disorder which is the collection of pathogens in the body. These pathological processes are occurrent parts of the *ogms:disease course* tuberculosis. If a patient has only lung tuberculosis and is completely healed from it without any sequelae, then the question is whether his instance of *ogms:disease course* is also an instance of *ogms:pathological process*. Under the current definitions of OGMS, and without axiomatization, it seems to speak to reason. If however taking a taxi to go and see a pneumologist would be qualified as one of the processes through which the patient’s disease is realized as well, then of course it is not.

There are 132 *sct:disorders* that have more than one role group and that are the domain of two distinct relationships with *sct:pathological process (attribute)*. In these cases, the *sct:disorder* being defined is not an instance of *ogms:pathological process*, but either of *ogms:disease course* or *ogms:disease*.

6. Conclusion

Both Schulz and we propose a solution towards compatibility of BFO/OGMS and SNOMED CT re diseases. Schulz brings SNOMED CT’s disease concepts under BFO’s occurrent hierarchy by reinterpreting their meaning and expressing the meaning through BFO-relations stripped from temporal indexing and restricted to what can be reasoned with in EL++. Our approach is to keep BFO’s axioms intact, and to create bridging axioms in FOL between BFO/OGMS and SNOMED CT. Schulz accepts the existence

of the various distinct disease-related types of entities that OGMS recognizes and uses these types in the motivation for his approach but not in his axioms. Our approach is to explicitly represent these types in the consequent parts of our bridging axioms. Our analysis indicates that it is possible to identify patterns in SNOMED CT's disorder concept definitions on the one hand, and definitions and axioms in BFO and OGMS on the other hand, that would make an automatic translation possible for a fair amount of disorders. A limitation of our approach is that there are many exceptions which require manual scrutiny. Further analysis might discover more opportunities for automation.

Acknowledgements

This work was supported in part by Clinical and Translational Science Award UL1 TR001412 from the National Institutes of Health. We thank the four anonymous reviewers for the detailed comments they provided.

References

- [1] Snomed International. Snomed CT 5-Step briefing.: Snomed International,; 2022 [cited 2023 January 6]. Available from: <https://www.snomed.org/snomed-ct/five-step-briefing>.
- [2] Rodrigues JM, Schulz S, Mizen B, Trombert B, Rector A. Scrutinizing SNOMED CT's Ability to Reconcile Clinical Language Ambiguities with an Ontology Representation. *Stud Health Technol Inform*. 2018;247:910-4. Epub 2018/04/22. PMID: 29678093.
- [3] Chang E, Mostafa J. The use of SNOMED CT, 2013-2020: a literature review. *J Am Med Inform Assoc*. 2021;28(9):2017-26. Epub 2021/06/22. doi: 10.1093/jamia/ocab084. PMID: 34151978; PMC8363812.
- [4] Bona JP, Ceusters W. Mismatches between major subhierarchies and semantic tags in SNOMED CT. *J Biomed Inform*. 2018;81:1-15. Epub 2018/02/21. doi: 10.1016/j.jbi.2018.02.009. PMID: 29462668.
- [5] Ceusters W, Mullin S. Expanding Evolutionary Terminology Auditing with Historic Formal and Linguistic Intentions: A Case Study in SNOMED CT. *Stud Health Technol Inform*. 2019;264:65-9. Epub 2019/08/24. doi: 10.3233/SHT1190184. PMID: 31437886.
- [6] Rossander A, Lindskold L, Ranerup A, Karlsson D. A State-of-the Art Review of SNOMED CT Terminology Binding and Recommendations for Practice and Research. *Methods Inf Med*. 2021;60(S 02):e76-e88. Epub 2021/09/29. doi: 10.1055/s-0041-1735167. PMID: 34583415; PMC8714300.
- [7] Smith SW, Koppel R. Healthcare information technology's relativity problems: a typology of how patients' physical reality, clinicians' mental models, and healthcare information technology differ. *J Am Med Inform Assoc*. 2014;21(1):117-31. Epub 2013/06/27. doi: 10.1136/amiajnl-2012-001419. PMID: 23800960; PMC3912703.
- [8] Rodrigues JM, Schulz S, Mizen B, Rector A, Serir S. Is the Application of SNOMED CT Concept Model sufficiently Quality Assured? *AMIA Annu Symp Proc*. 2017;2017:1488-97. Epub 2018/06/02. PMID: 29854218; PMC5977666.
- [9] Schulz S, Rodrigues JM, Rector A, Spackman K, Campbell J, Ustun B, et al. What's in a class? Lessons learnt from the ICD - SNOMED CT harmonisation. *Stud Health Technol Inform*. 2014;205:1038-42. Epub 2014/08/28. PMID: 25160346.
- [10] Schulz S, Rector A, Rodrigues JM, Spackman K. Competing interpretations of disorder codes in SNOMED CT and ICD. *AMIA Annu Symp Proc*. 2012;2012:819-27. PMID: 23304356; 3540515.
- [11] Schulz S, Cornet R. SNOMED CT's Ontological Commitment. In: Smith B, editor. *ICBO: International Conference on Biomedical Ontology*. Buffalo NY: National Center for Ontological Research; 2009. p. 55-8.
- [12] Scheuermann RH, Ceusters W, Smith B. Toward an ontological treatment of disease and diagnosis. *Summit Transl Bioinform*. 2009;2009:116-20. Epub 2009/01/01. PMID: 21347182; PMC3041577.
- [13] Smith B, Ceusters W. Ontological realism: A methodology for coordinated evolution of scientific ontologies. *Appl Ontol*. 2010;5(3-4):139-88. Epub 2011/06/04. doi: 10.3233/AO-2010-0079. PMID: 21637730; PMC3104413.
- [14] Hogan W. Aligning the top-level of SNOMED-CT with Basic Formal Ontology. *Nature Precedings*. 2008. doi: 10.1038/npre.2008.2373.1.

- [15] El-Sappagh S, Franda F, Ali F, Kwak KS. SNOMED CT standard ontology based on the ontology for general medical science. *BMC Med Inform Decis Mak.* 2018;18(1):76. Epub 2018/09/02. doi: 10.1186/s12911-018-0651-5. PMID: 30170591; PMC6119323.
- [16] Schulz S. SNOMED CT x BFO: can the gap between legacy terminology and foundational ontology be bridged? ICBO/FOIS; September 16, 2021; Bolzano, Italy2021.
- [17] Pengput A, Ceusters W. Setting the Scene to Link SNOMED CT to Realism-Based Ontologies 19th World Congress on Medical and Health Informatics (MEDINFO); 8 - 12 July 2023; Sydney, Australia.
- [18] Snomed International. Snomed International SNOMED CT browser: Snomed International.; 2023 [cited 2023 January 6,]. Available from: <https://browser.ihtsdotools.org/>
- [19] International Health Terminology Standards Development Organisation. SNOMED CT® Editorial Guide - v20230131. 2023. <https://confluence.ihtsdotools.org/display/DOCEG/PDFs+for+Download>.
- [20] Schulz S, Romacker M, Hahn U. Part-whole reasoning in medical ontologies revisited--introducing SEP triplets into classification-based description logics. *Proceedings / AMIA Annual Symposium AMIA Symposium.* 1998:830-4. PMID: 9929335; 2232338.
- [21] Spackman KA, Dionne R, Mays E, Weis J. Role grouping as an extension to the description logic of Ontylog, motivated by concept modeling in SNOMED. *Proceedings / AMIA Annual Symposium AMIA Symposium.* 2002:712-6. PMID: 12463917; 2244464.
- [22] Klein GO, Smith B. Concept Systems and Ontologies: Recommendations for Basic Terminology. *Trans Jpn Soc Artif Intell.* 2010;25(3):433-41. Epub 2011/02/11. doi: 10.1527/tjsai.25.433. PMID: 21308002; PMC3034144.
- [23] Ceusters W, Smith, B. Ontology and Medical Terminology: Why Description Logics are not enough. *Proceedings of the Conference Towards an Electronic Patient Record (TEPR 2003).* 2003.
- [24] International Standards Organisation. ISO/IEC 21838-2:2021 Information technology — Top-level ontologies (TLO) — Part 2: Basic Formal Ontology (BFO). 2021.
- [25] Flügel S, Glauer M, Neuhaus F, Hastings J. When one Logic is Not Enough: Integrating First-order Annotations in OWL Ontologies. *Semantic Web Journal [Internet].* 2022. Available from: <https://arxiv.org/abs/2210.03497>.
- [26] Ceusters W, Blaisure J. Caveats for the Use of the Active Problem List as Ground Truth for Decision Support. *Stud Health Technol Inform.* 2018;255:10-4. Epub 2018/10/12. PMID: 30306897.
- [27] Ceusters W. The Place of Referent Tracking in Biomedical Informatics. In: Elkin PL, editor. *Terminology, Ontology and their Implementations : Teaching Guide and Notes.* Cham: Springer International Publishing; 2022. p. 39-46.
- [28] Ceusters W. Mutualistic integration of Referent Tracking and BFO2020-style FOL axiomatization Buffalo, NY2022 [updated October 24, 2022.]. Available from: <http://www.referent-tracking.com/RTU/files/RTandBFO2020-20221024.pptx>.
- [29] Smith B, Ceusters W. Aboutness: Towards Foundations for the Information Artifact Ontology. *International Conference on Biomedical Ontology; July 27-30; Lisbon, Portugal 2015.* p. 47-51.