# FHIR Profiles and Consolidated CDA Templates: Proposed Data-Modeling Changes With Benefits for Patient Safety

## Future State

July 31, 2018

*DRAFT*

**Contents**

# 1 Introduction

A close inspection of commonly used FHIR profiles and C-CDA templates reveals a number of modeling issues that present potential safety risks when these standards are used to transmit real-world patient data. A common theme across these modeling issues is "under-specification," i.e. the defined standards allow senders of important clinical data to represent the information in a manner that is not expected or not understood by receivers, although both parties are nominally conformant to the standard specifications. A prior report detailed these issues[1] and the ways in which they can compromise patient safety (in the text below, this prior report is referred to as the "Current State Whitepaper"). The present report ("Future State Whitepaper") proposes strategies to address the identified issues in future versions of the profiles and implementation guides.

The first part of the report proposes solutions based on the definition of more granular and specific models of clinical data than those currently specified in FHIR and C-CDA standards. The second part discusses a more formal modeling of negation to ensure that meanings of negated clinical findings are precise and correctly interpreted by data recipients. The third part addresses several specific issues related to the specification of controlled terminologies and value sets. The final part proposes several miscellaneous modeling improvements to improve the clarity and semantic consistency of the standards.

# 2 Introducing More Detailed Clinical Models

A key reason that under-specification still exists within FHIR profiles and C-CDA templates is that existing profiles and templates remain quite generic, i.e. they are intended to represent a large number of potential types of clinical observations, conditions, procedures, etc. Their designs must, therefore, be general enough to accommodate various ways of representing clinical data, because different specific types of clinical data have different modeling requirements. For example, although hematology results, microbial culture results, pulmonary function results, and physical exam findings are all instances of the FHIR *Observation* profile, they each require different modeling to represent the relevant clinical information.

Further, data elements relevant to only certain of these types of clinical observation cannot be reasonably included as optional elements of generic profiles and templates, because the number of such observation-specific data elements is large and would bloat and obfuscate the generic profiles and templates that currently exist. Lastly, the value sets for coded data elements must currently be large enough to accommodate any of the types of clinical data to be represented, and, even then, adherence to the value sets needs to be optional to accommodate unexpected types of clinical data that generic profiles and templates need to cover.

The natural solution to these issues is the definition of more detailed profiles and templates that are specifically designed to be used to represent only certain types of clinical data. For example, rather than having a generic *Observation* profile, an implementation guide should have hundreds of more specific profiles for different types of observations that have different modeling requirements. Initiatives are underway to define such "detailed clinical models" (DCMs), such as the "Clinical Information Modeling Initiative" (CIMI) at HL7[2]. As described below, such initiatives have the potential to remedy certain of the patient-safety issues inherent in the generic design of FHIR profiles and C-CDA templates that prevail today.

## 2.1 Benefit: Standardizing Structure with Respect to Predefined Data Elements

As discussed in Section 2.2.1 of the Current State Whitepaper, the FHIR US-Core implementation guide contains profiles that remain underspecified with respect to the representation of specific types of clinical observations. Specifically, the profiles allow the relevant components of a clinical observation to be structured in multiple different ways, each equally conformant with the constraints of the FHIR implementation guide. For example, as illustrated in the Current State Whitepaper, an observation as simple as a blood-pressure measurement can be represented structurally in at least three different ways, i.e., using varying elements of FHIR profiles and varying numbers of linked FHIR profile instances.

The specification of more detailed FHIR profiles that prescribe the structures to use for various types of specific clinical observations (such as blood pressure, culture results, pulmonary function tests, etc.) can address this problem. In fact, such detailed profiles were specified in the FHIR US-Core implementation guide for a small number of clinical observations, primarily vital-sign measurements (see Figure 1, showing how HL7 selected method "C" among the three options for representing blood pressure readings that were shown in Section 2.2.1 of the Current State Whitepaper).

The definition of a large number of detailed profiles was beyond the resource and time constraints of the Argonaut and HL7 groups that defined the US-Core implementation guide. However, many more detailed profiles within the FHIR US-Core Implementation Guide are required for a much broader set of clinical observations, including specific kinds of lab test results (chemistry versus culture), functional test results (pulmonary function tests versus Apgar assessments), and clinical exam findings (cardiac exam versus neurological exam).

At the same time, many such detailed profiles have been defined for one-off uses by isolated groups (see, for example, the library of FHIR profiles posted at Simplifier.Net[3]. However, these libraries of profiles are neither canonical (i.e., they can and do contain multiple, alternative profiles for the same clinical observations), nor authoritative (i.e., the profiles they contain are not part of broader, coherent implementation guides, such as FHIR US-Core, to which interoperating organizations are obliged to conform). Hence, considerable additional work is needed to evaluate such FHIR profiles, select the best ones for each clinical observation of interest, and designate those as formal profiles within implementation guides such as FHIR US-Core.

**Figure 1. The detailed HL7 FHIR profile for the clinical observation of "blood pressure".**

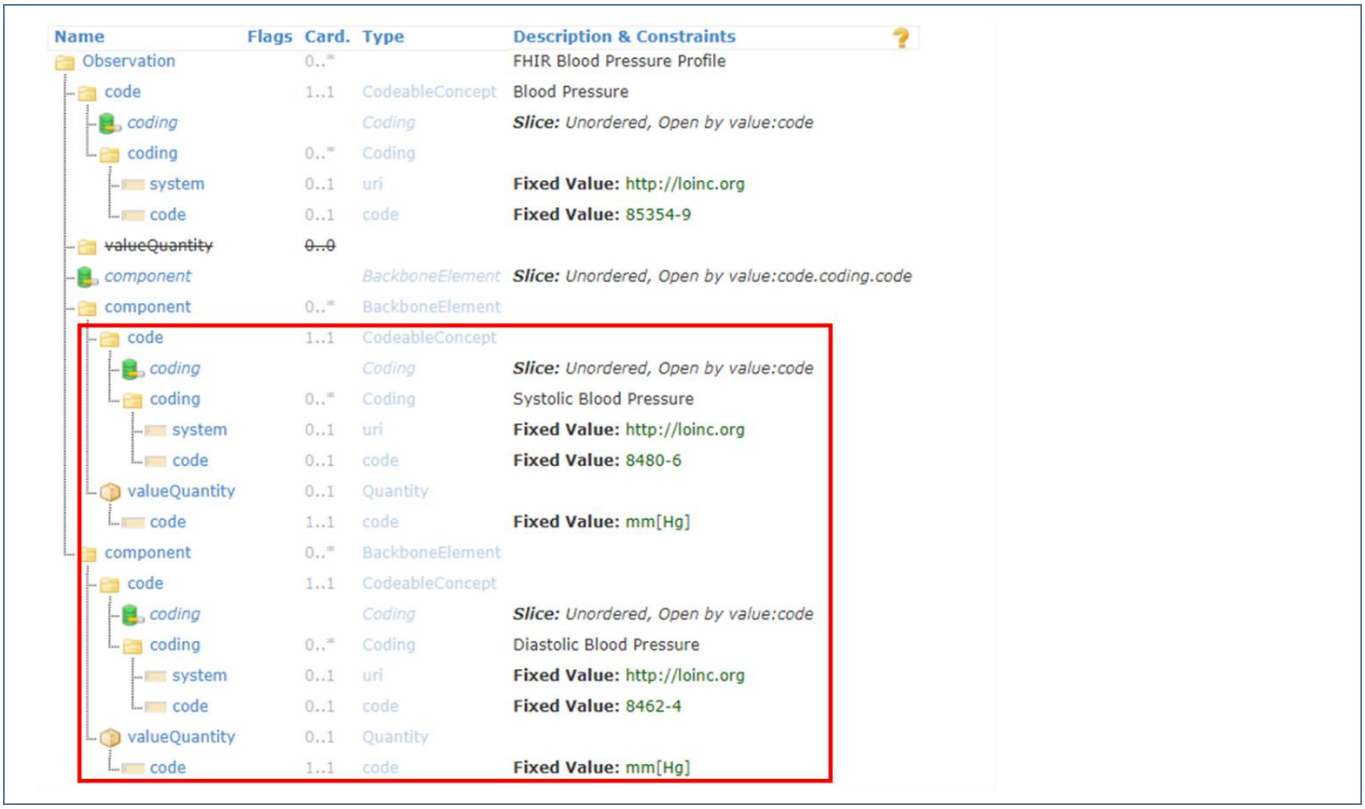


The DCMs developed by the CIMI project can serve as sources of domain content for such an undertaking, as CIMI is already specifying representational models for a large number of clinical observations, and these models are much more detailed than the existing profiles in the FHIR US-Core Implementation Guide[4]. By converting certain of these models to FHIR profile specifications, designating those profiles as formal components of the FHIR US-Core Implementation Guide, and ensuring that no alternative profiles are introduced into that implementation guide to (differently) represent the same clinical observations, HL7 could address the problem of structural underspecification that currently poses risks to patient safety among organizations using the FHIR US-Core implementation guide to exchange important clinical data.

## 2.2 Benefit: Standardizing Structure with Respect to Needed Extension Elements

As discussed in Section 2.2.2 of the Current State Whitepaper, the profiles of the FHIR US-Core Implementation Guide include only a core set of modeling elements intended to represent the most common aspects of clinical observations. Detailed aspects of particular types of observations, such as a patient's body position when blood pressure is measured, are omitted from the profiles. If the sender of a FHIR profile instance wishes to include such information within a blood pressure measurement, it must add an "extension" element to the profile instance. However, such elements need not be (and generally are not) defined within the FHIR US-Core Implementation Guide, and therefore may be entirely unexpected and unrecognized by recipients who are using only that implementation guide as the basis for their FHIR-processing implementations.

The specification of more detailed FHIR profiles, as discussed in Section 2.1, however, creates the opportunity to address this issue. Specifically, the existence of more detailed FHIR profiles for clinical observations, such as blood pressure measurements, culture results, and pulmonary function tests,

creates the opportunity to formally define and add specific extension elements to such detailed FHIR profiles only. Therefore, the extension element of "body-position," for example, could be added to the FHIR profile for blood pressure measurements. That profile would be made a part of the FHIR US-Core Implementation Guide that is referenced by all implementers of FHIR-generating and FHIR-consuming processes. The result would be that all senders of FHIR profile instances containing a patient's blood pressure would be obligated to represent the patient's body position using the same extension element, and the structure and coding of that extension element would be known to any recipient of the FHIR profile instance.

For example, Figure 2 shows the FHIR profile for "blood pressure" as defined by HL7 with the envisioned addition of a specific pre-defined extension element for "Body Position."

**Figure 2. The FHIR profile for "blood pressure" with a defined extension for "Body Position"**

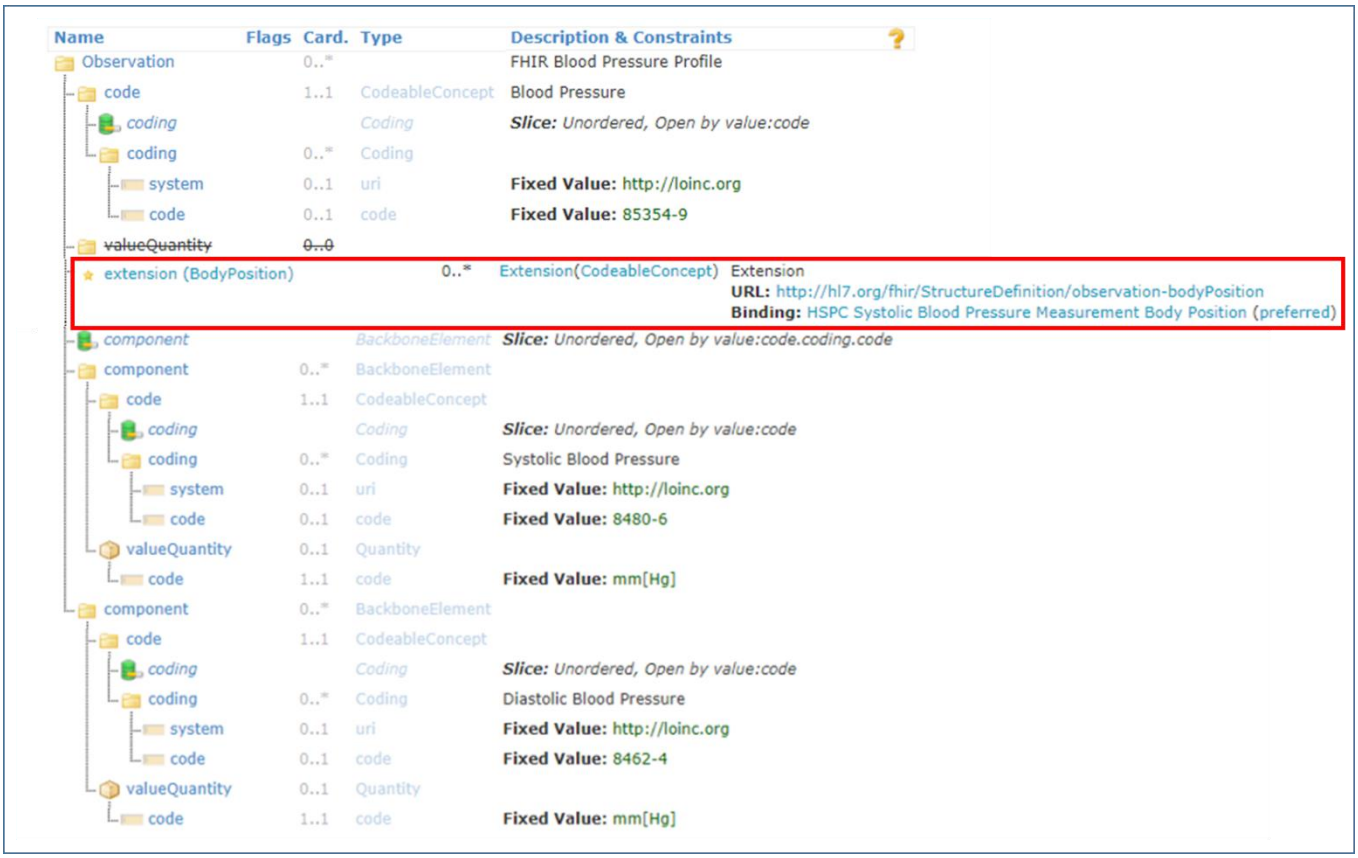


## 2.3 Benefit: Further Constraints on Value Sets

As discussed in Sections 2.6.2 and 3.5.2 of the Current State Whitepaper, certain value sets specified for data elements in both the FHIR US-Core Implementation Guide and the C-CDA Implementation Guide specify a "SHOULD" binding constraint, rather than a "SHALL" binding constraint. A "SHOULD" constraint means that conformant sending systems are *encouraged* to use the specified value set when populating such data elements, but are not obligated to, i.e., data instances that do not use the recommended value sets are still conformant to the implementation guides and will pass formal validation testing. The patient-safety ramifications of this approach are that the recipients of FHIR resource instances or C-CDA template instances that include codes outside of recognized value sets may

be unable to automatically recognize and use such codes within decision-support functions, possibly "breaking" decision support rules and leading to false-negative clinical alerts or reminders.

The development of more specific FHIR profiles for particular types of clinical observations can also address this current limitation of the FHIR US-Core and C-CDA Implementation Guides. Specifically, in the context of particular clinical observations, it may be possible to categorically require codes from more constrained value sets that are specific to the clinical context of those observations. For example, the FHIR US-Core *Observation* resource profile specifies that implementers "SHOULD" use only codes from SNOMED-CT when populating the "value" data element in instances of this profile (but are allowed to use other codes). If a more specific profile existed just for culture results, in which the coded values were intended to represent only the identities of cultured organisms, it would be possible to enumerate such potential organisms comprehensively within a more constrained value set. The implementation guide could thereby require sending parties to always populate the value element in culture results with only one of these codes (i.e., designate the value-set binding as a "SHALL" constraint, rather than a "SHOULD" constraint), rather than allow sending parties to, effectively, decide to use other coding systems at their discretion (as the current implementation guide allows).

# 3 Formalizing Negation in FHIR and C-CDA

A second type of underspecification in FHIR US-Core profiles and C-CDA templates relates to the manner(s) in which clinical statements are negated. Expressing the negation of clinical findings, diagnoses, treatments, etc. is an important aspect of clinical documentation, and an important input to data-analysis and decision-support processes. For example, it may be equally significant to know that a patient has never had a history of peptic ulcer disease or a family history of atherosclerosis in when diagnosing a current ailment as it is to know the affirmative presence of such a history. As described below, current mechanisms for expressing negation in FHIR profiles and C-CDA templates could and should be improved to better express such negated statements when exchanging clinical data.

## 3.1 Denoting Scope of Negation

As discussed in Sections 2.3.2 and 3.4.2 of the Current State Whitepaper, current methods of negating clinical statements in FHIR profiles and C-CDA templates do not require or even allow the specification of the *scope* of a negation. Instances of these profiles and templates include myriad clinical data, context data, and meta-data that collectively represent the entire clinical observation. The scope of negation denotes the specific components of a clinical observation that are intended to be negated. Importantly, the precise semantics of a negated clinical observation depend on formally specifying the set of such components that are included in the scope of negation.

For example, a clinical statement may express that 50 units of pneumovax vaccine from lot 14873 were administered to a patient on July 10, 2018, and the administration was recorded by Nurse Morris. The negation of the entirety of the clinical statement, however, would formally express that only the *entire conjunction* of the expressed statements did not occur. However, it would not negate, for example, the administration of 75 units of pneumovax or the administration of 50 units of pneumovax recorded by Nurse Smith rather than Nurse Morris instead. If the intent of the negation is to express that *no* administration of pneumovax took place whatsoever on the indicated date, then the scope of negation must be explicitly denoted to represent that.

One method of representing the scope of negation for nested hierarchical structures, such as FHIR JSON resources or C-CDA XML documents, is via the use of path expressions. Path expressions denoted in XPath[5] or JSONPath[6] identify the specific elements and attributes within a clinical statement that are intended to be negated (i.e., the conjunction of which is logically negated). The explicit representation

of this information allows automated algorithms to correctly process negated clinical statements in the course of making inferences, including inferences important to data analysis and clinical decision support.

For example, Figure 3 shows an instance of an Immunization Activity template as defined by the C-CDA Implementation Guide with an additional (envisioned) element "negationScope" added. NegationScope is intended to denote the scope of negation for the instance by specifying the specific expressions within the XML hierarchy whose conjunction is negated. Note that only a subset of the attributes and elements nested within the Immunization Activity element have been negated in this example, specifically the mood code, the effective time, and identity of the immunization. The formal semantics of the negated clinical observation in Figure 3, therefore, is that no event occurred on July 10, 2018 in which a pneumovax vaccine was administered to the patient.

Importantly, other components of the clinical statement are not negated, such as the reason the vaccine was not administered and the person who documented the non-administration of the vaccine. Hence, those aspects of the clinical statement remain affirmatively true (as they should in this case). Further, by excluding from the scope of negation other details of the vaccination event that are expressed in the negated clinical statement (such as the route, dose, and lot number of the vaccine), the semantics of the negation are expanded to include the administration of pneumovax *in any dose by any route from any lot number* on July 10, 2018, which is likely the intended meaning of the negated clinical statement. Conversely, for example, if the lot number were included in the scope of negation, a decision-support algorithm could not logically exclude the possibility that a pneumovax vaccine from a different lot number had been administered to the patient on July 10, 2018[1].

[1] Note that the negation of a conjunction (such as "NOT(pneumovax AND 2018-07-10 AND Lot-14873)") is logically equivalent to the *disjunction* of its negated constituents (i.e., "NOT(pneumovax) OR NOT(2018-07-10) OR NOT(Lot-14873)") rather than the conjunction of its negated constituents (i.e., "NOT(pneumovax) AND NOT(2018-07-10) AND NOT(Lot-14873)"). Disjunctions and conjunctions of negated conditions are logically different because, more generally, "NOT(A) OR NOT (B)" is true if *either* A or B is false, whereas "NOT (A) and NOT(B)" is true only if *both* A and B are false. Hence, any query expression that did not specify a lot number would be satisfied by the negation of such a conjunction (e.g., "IF [ (pneumovax) AND (2018-07-10) ] THEN…"), which is likely not the intent of the negated clinical statement and could cause decision-support rules that include such a query to return incorrect results if the scope of negation in the clinical statement were specified incorrectly.

**Figure 3. Example of a negated clinical statement in which the scope of negation is explicitly denoted.**

```
<substanceAdministration classCode="SBADM" moodCode="EVN" negationInd="true">  <!-- ** Immunization Activity template** -->
  ...lines omitted...
  <negationScope>
    <entry path="../../@moodCode"/>
    <entry path="../../effectiveTime"/>
    <entry path="../../consummable/manufacturedProduct/manufacturedMaterial/code"/>
  </negationScope>
  <effectiveTime value="20180710"/>
  <routeCode code="C28161" codeSystem="2.16.840.1.113883.3.26.1.1" displayName="Intramuscular injection"/>
  <doseQuantity value="50" unit="ug"/>
  <consumable>
    <manufacturedProduct classCode="MANU">   <!-- ** Immunization Medication template ** -->
      ...lines omitted...
      <manufacturedMaterial>
        <code code="33" codeSystem="2.16.840.1.113883.6.59" displayName="Pneumococcal polysaccharide vaccine" codeSystemName="CVX"/>
        <lotNumberText>14873</lotNumberText>
      </manufacturedMaterial>
    </manufacturedProduct>
  </consumable>
  <entryRelationship typeCode="RSON">
    <observation classCode="OBS" moodCode="EVN">  <!-- ** Immunization Refusal Reason template ** -->
      <templateId root="2.16.840.1.113883.10.20.22.4.53"/>
      <id root="2a620155-9d11-439e-92b3-5d9815ff4dd8"/>
      <code displayName="Patient Objection" code="PATOBJ" codeSystem="2.16.840.1.113883.5.8"/>
      <statusCode code="completed"/>
    </observation>
  </entryRelationship>
  <author>  <!-- ** Author Participation template ** -->
    ...lines omitted...
    <assignedAuthor>
      <id root="20cf14fb-b65c-4c8c-a54d-b0cca834c18c"/>
      <assignedPerson>
        <name>
          <given>Gail</given><family>Morris</family><suffix>RN</suffix>
        </name>
      </assignedPerson>
    </assignedAuthor>
  </author>
</substanceAdministration>
```

## 3.2 Adding Negation Attributes Where Needed

As discussed in Sections 2.3.3 and 3.4.3 of the current whitepaper, there exist certain FHIR resource and profile definitions, as well as C-CDA template specifications, that omit the ability to negate applicable clinical statements. For example, the FHIR *Observation* resource and the C-CDA *Result Observation* template both lack data elements to explicitly negate observed clinical findings.

In the absence of standardized mechanisms to express negation within these data structures, the senders of FHIR resource instances and C-CDA template instances must resort to *ad hoc* negation methods when they need to express negated clinical statements. Such methods include the use of coded values that express negation in a pre-coordinated fashion for certain findings (such as SNOMED-CT code 301131000 – "Heart murmur absent"), the use of Boolean-typed values for coded observation types (such as observation type = "Heart murmur quality", value = "false"), or the use of numeric-typed values for coded observations (such as observation type = "Heart murmur grade", value = 0). However, *Ad hoc* representations will tend to vary across different senders of clinical data, and recipients may be unaware of the specific techniques or (or kludges) that any given sender may employ.

A preferred method is to include explicit negation attributes in all clinical statement models when negated data are likely to be recorded and to be important for data analysis or clinical decision support.

Such attributes could comprise Boolean negation flags, such as the existing "negationInd" attribute available in a number of C-CDA templates. Note that "actionNegationInd"[7] and "valueNegationInd"[8] are negating attributes defined within the HL7 V3 Reference Information Model (RIM) for any clinical "Act" or "Observation," respectively. As both the FHIR standard and the C-CDA standard are based on the RIM, these attributes would be natural additions to any clinical data structures based on the RIM Act and Observation classes.

Our current analysis has identified the following data structures within the FHIR and C-CDA standards that would benefit from having such explicit negating attributes added:

- FHIR
  - Observation (base FHIR resource)
  - FamilyMemberHistory (US-Core FHIR resource profile)
  - Procedure (US-Core FHIR resource profile)

- C-CDA
  - Result Observation Template
  - Family History Observation

# 4 Improving Terminology and Value-Set Restrictions

A third area in which changes and improvements to the modeling of clinical data are required with the FHIR US-Core and C-CDA Implementation Guides is in the management of terminology constraints and value-set constraints.

## 4.1 Disallow use of SNOMED-CT Codes from "Situation-with-Explicit-Context" Hierarchy

A number of important data elements in FHIR profiles and C-CDA templates specify coded value sets that allow codes from the SNOMED-CT "Situation-with-Explicit-Context" Hierarchy (i.e., codes from the SNOMED-CT sub-tree rooted at the code 243796009). This SNOMED hierarchy contains pre-coordinated codes to represent findings that are formally annotated with contextual qualifiers, such as "Finding Context" (to denote whether a finding is present, absent, or uncertain), "Temporal Context" (to denote whether the finding is currently present or was present in the past), and "Subject Relationship Context" (to denote whether a finding applies to a patient or a family member of a patient)[9]. Figure 4 shows the graphical rendition of such a code, in this case representing the concept "No family history of hypertension."

**Figure 4. Example of a SNOMED-CT code from the "Situation-with-Explicit-Context" hierarchy.**

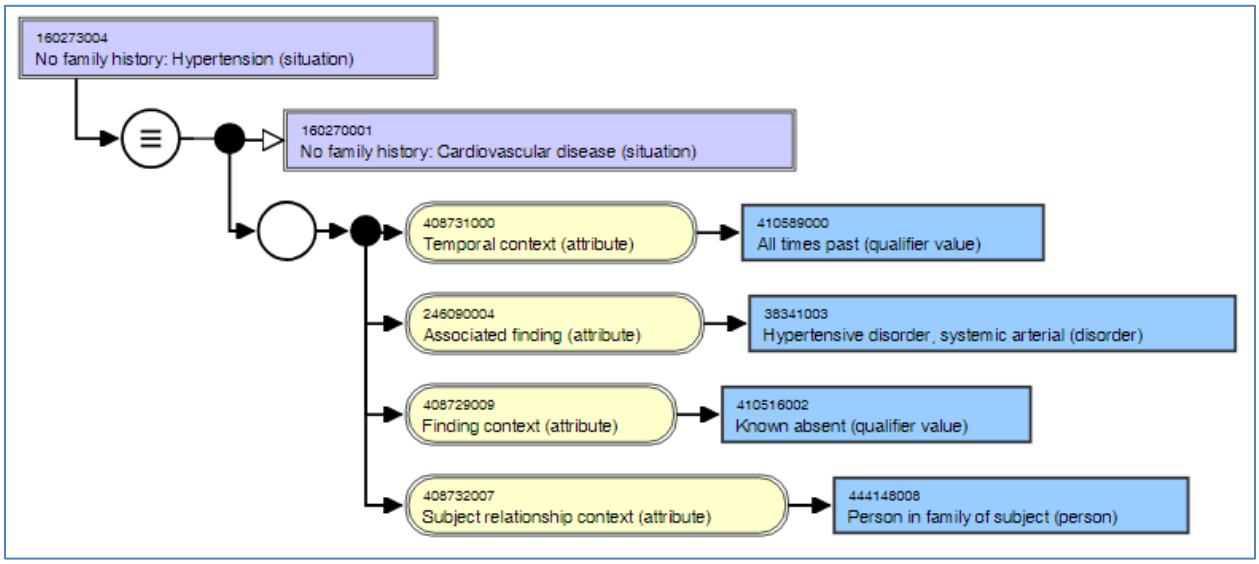


Examples of data elements that may be populated with SNOMED "situation-with-explicit-context" codes in FHIR profiles and C-CDA templates include the "code" field used to denote patient diagnoses in the FHIR US-Core *Condition* profile and the "value" field used to denote patient diagnoses in the C-CDA *Problem Observation* template. As discussed in Sections 2.6.1 and 3.4.1 of the Current State Whitepaper, however, including "situation-with-explicit-context" codes in the value sets for these data elements introduces at least two problems.

1. **Redundancy in representing negated findings**. The FHIR US-Core *Condition* profile and the C-CDA *Problem Observation* template both also include data elements for explicitly negating clinical statements. Hence, a clinical statement representing the absence of a finding could be represented in two distinct ways: (1) specifying a "situation-with-explicit-context" code from SNOMED-CT that directly negates the finding (as shown in Figure 4, above), *or* (2) specifying the affirmative clinical finding (as shown in Figure 5, below), but negating the entire clinical statement using a separate data element, such as "verificationStatus" in the FHIR US-Core *Condition* profile or "negationInd" in the C-CDA *Problem* Observation template. The availability of two redundant representations for the same clinical statement increases the risk that one of the representations will not be used or recognized by a recipient of the clinical information.

**Figure 5. The affirmative finding referenced in the pre-coordinated code in Figure 4.**

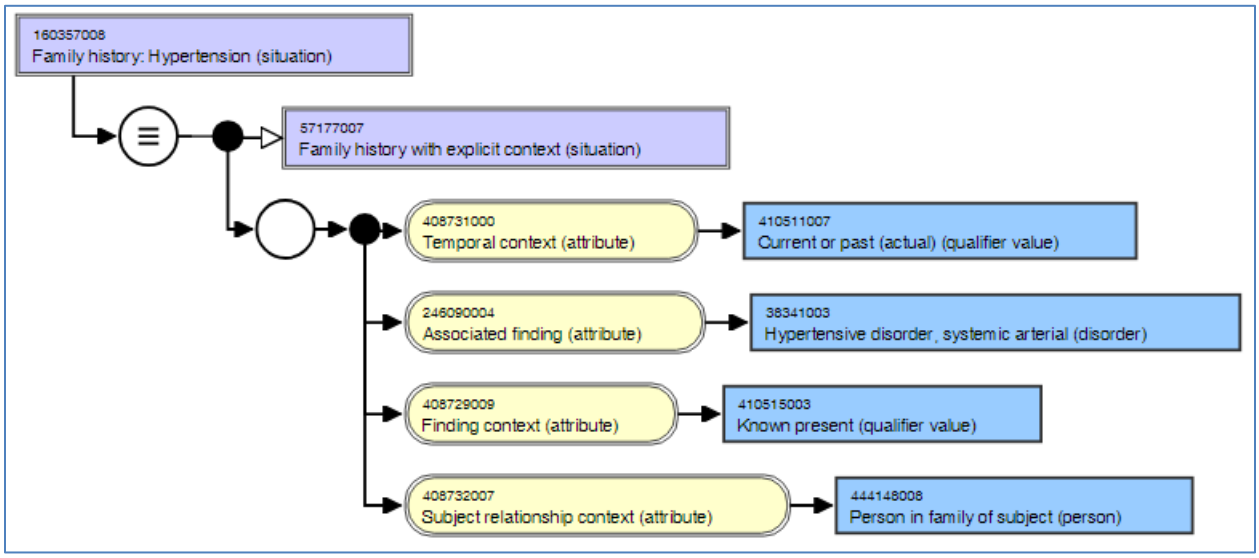


2. **Incorrect subsumption-testing results for negated findings.** Logical subsumption testing may not work correctly over SNOMED-CT codes from the "situation-with-explicit-context" hierarchy when such codes include "Finding Context" values of "Known absent" (i.e., when such codes represent negated clinical statements). For example, consider the SNOMED-CT codes in Figure 4 ("No family history of hypertension") and Figure 6 ("No family history of heart disease). Given the definitions of these concepts in SNOMED-CT and the rules of logical subsumption, "No family history of heart disease" will subsume "No family history of hypertension" because the finding "heart disease" subsumes the finding "hypertension". But this subsumption inference is patently incorrect. Based on this behavior, any patient with a recorded finding of "No family history of hypertension" will satisfy a query for patients with "No family history of heart disease." However, such patients could have family histories of angina or atherosclerosis or any number of other heart diseases, and hence should not satisfy such a query. The error occurs because negation does not follow the normal logical rules of subsumption.[2]

[2] In fact, the opposite subsumption-testing behavior occurs with negated findings. Specifically, if "heart disease" subsumes "hypertension", then "no family history of hypertension" subsumes "no family history of heart disease", because a patient known to have no family history of any heart disease could not have a family history of hypertension.

**Figure 6. Example of a subsuming SNOMED-CT code from the "Situation-with-Explicit-Context" hierarchy.**

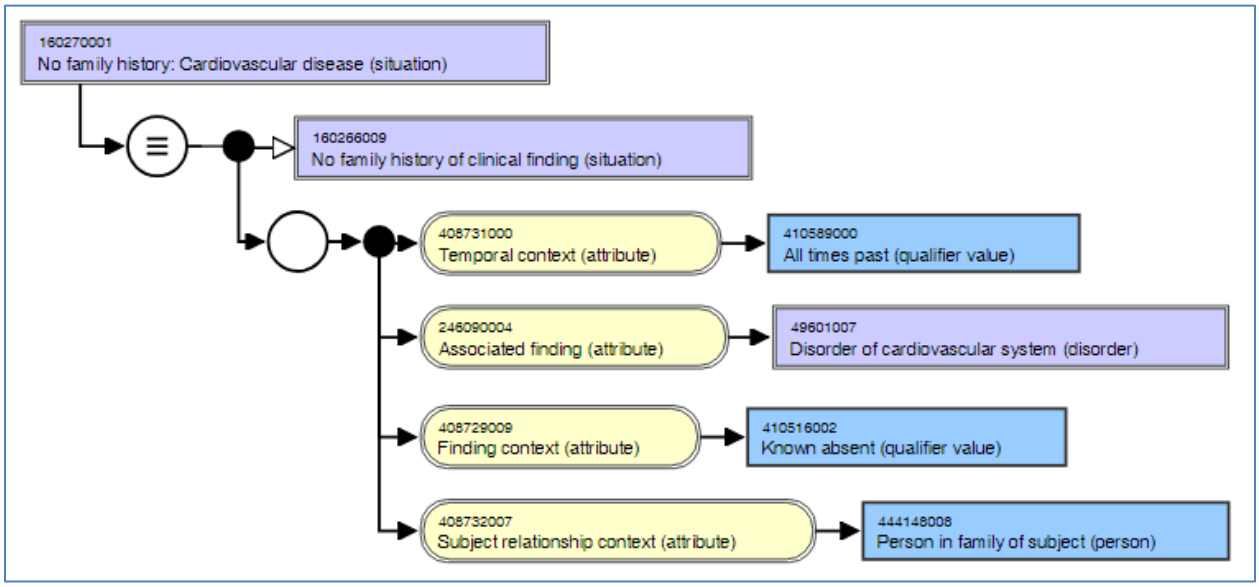


Hence, codes from the SNOMED-CT "situation-with-explicit-context" hierarchy should not be allowed in FHIR US-Core profile instances and C-CDA templates, because such codes may result in transmitted clinical findings not being recognized by receiving entities, as well as not being correctly handled by subsumption-testing algorithms.

## 4.2 Use of Terminology Consolidation via SOLOR for Multi-Terminology Value Sets

As discussed in Sections 2.6.1 and 3.5.1 of the Current State Whitepaper, certain value sets specified in FHIR US-Core profiles and C-CDA templates include overlapping codes from multiple coding systems. For example, the FHIR US-Core *AllergyIntolerance* profile and the C-CDA *Allergy Intolerance Observation* template both specify a value set for denoting allergenic substances that includes codes from NDF-RT, RxNorm, UNII, and SNOMED-CT.

However, this approach results in a value set with multiple codes for the same concepts. For example, RxNorm, UNII, and SNOMED-CT all contain a code for the drug ingredient "Metronidazole hydrochloride" (RXCUI 82047, UNII 76JC1633UF, and SCTID 3941003), as well as many other drug ingredients. Although both FHIR and C-CDA specify a prioritization rubric that prescribes the appropriate coding system for different types of allergens (e.g., drug classes versus drug ingredients)[10], it is not unlikely that the senders of allergy intolerance information may stray from this rubric (accidentally or for convenience) and use a different coding system than prescribed. In such cases, validation engines for FHIR US-Core profiles and C-CDA templates will not detect the errors because codes from all of the coding systems are technically allowed. Receiving systems that do not expect or do not recognize the transmitted codes in such erroneous situations may misinterpret important drug-allergy information.

To mitigate such situations, it is important for senders and recipients to use a consolidated terminology resources that rigorously represents, classifies, and maps coded concepts from disparate coding systems. One such resource is the SOLOR terminology system[11], which integrates specific terminologies necessary for clinical data representation, such as SNOMED-CT, LOINC, RxNorm, and others. Using the SOLOR Common Model and existing terminology content, one could create an integrated terminology

resource that includes relevant codes from NDF-RT, RxNorm, UNII, and SNOMED-CT and represents all of the synonymous and hierarchical relationships among them.

Such a resource could, for example, automatically include the correctly prioritized codes in FHIR profile instances and C-CDA templates that senders of data generate, regardless of which coding system the senders used internally to represent allergenic substances. Alternatively, such a resource could be used by recipients of FHIR profile instances and C-CDA template instances to correctly map whatever codes they receive in FHIR profile instances and C-CDA template instances to the appropriate codes they were expecting based on the prioritization rubric specified in these standards.

# 5 Additional Miscellaneous Recommendations

## 5.1 Better Specify/Formalize use of SCT Post-Coordinated Expressions as Codes

As discussed in Section 3.5.3 of the Current State Whitepaper, the C-CDA *Problem Observation* template specifies that the coded values of problems may include optional qualifiers that further modify the stated values. For example, as shown in Figure 7, a qualifier may specify the laterality of an ulnar fracture if no pre-coordinated code exists for "Fracture of the left ulna," but that is the clinical observation that a sender wishes to communicate.

**Figure 7. Example of a code qualifier in a C-CDA *Problem Observation* template instance.**

```
<value code=" 54556006" codeSystem="2.16.840.1.113883.6.96" displayName="Fracture of ulna">
  <qualifier>
    <name code="78615007" codeSystem="2.16.840.1.113883.6.96" displayName="with laterality"/>
    <value code="7771000" codeSystem="2.16.840.1.113883.6.96" displayName="left"/>
  </qualifier>
</targetSiteCode>
```

This useful capability to qualify codes is a feature of the underlying HL7 V3 RIM data model, which the C-CDA implementation guide also supports[12]. The feature is particularly useful when SNOMED-CT codes are used to represent coded problems (as suggested, but not required, by the C-CDA implementation guide), because SNOMED-CT defines a formal model for combining codes and modifiers into post-coordinated expressions, which substantially expands the expressive capability of SNOMED-CT.

However, the C-CDA implementation guide does not explicitly specify how to use the features of the HL7 RIM data model to represent post-coordinated SNOMED-CT expressions for values that may be coded using SNOMED-CT. Further, the implementation guide does not refer readers to documentation of the SNOMED-CT content model, which defines and constrains the specific post-coordinated SNOMED-CT expressions that may be expressed. In certain cases, the allowed expressions may need to be constrained further than SNOMED-CT allows to prevent the redundant representation of qualifiers already represented among the C-CDA template's data elements (such as the status of a problem). Finally, the implementation guide does not provide examples of post-coordinated SNOMED-CT expressions represented using the C-CDA XML schema (as shown in Figure 7). Future versions of the C-CDA implementation guide should include this additional documentation to guide users in the creation and consumption of SNOMED-CT post-coordinated expressions within C-CDA template instances, where allowed.

## 5.2 Prohibit Modifying Extension Elements in FHIR Instances that are not Defined in the FHIR Implementation Guide

As discussed in Section 2.4.2 of the Current State Whitepaper, FHIR resources and FHIR profiles (including the US-Core profiles) allow the senders of FHIR resource instances to include arbitrary additional data elements that are not included in the published definitions of the resource or profile, but can change the fundamental meaning of the entire resource instance. These so-called "modifying extension" data elements may present significant patient-safety risks when they appear in resource instances that are received by providers and that are then processed by automated decision-support systems. Because the modifying extension data elements are not specified in any published specifications for the received resources (i.e., in the FHIR US-Core Implementation Guide), decision support systems may not be programmed to recognize that these elements are changing the meaning of the received data. The result may be an incorrect or absent decision-support alert or recommendation which compromises patient safety.

For example, a sending application could theoretically add a modifier extension element to a *MedicationRequest* resource instance that indicates the patient should NOT take the medication as prescribed for the next 3 days. The FHIR standard allows the sending application to include such an extension element even if the element is not specified in the *MedicationRequest* resource or in any documented profile of this resource. A receiving application, therefore, may not expect or recognize the modifier extension, and therefore could incorrectly assume that the patient *should* take the medication as prescribed for the next 3 days, a potential patient-safety error.

For this reason, the FHIR US-Core implementation guide should be modified to exclude any modifying data elements that are not explicitly defined in resource profiles.

# 6 References

[1] FHIR Profiles and Consolidated CDA Templates: Data-Modeling Issues With Implications for Patient Safety. VA Knowledge Based Systems Group. April 6, 2018.

[2] http://www.hl7.org/Special/Committees/cimi/index.cfm. (Accessed 7/30/2018)

[3] https://simplifier.net/ui/search/index?category=Profile. (Accessed 7/29/2018)

[4] http://www.clinicalelement.com/cimi-browser/#/. (Accessed 7/30/2018)

[5] https://en.wikipedia.org/wiki/XPath. (Accessed 7/30/2018)

[6] http://goessner.net/articles/JsonPath/. (Accessed 7/30/2018)

[7] http://hl7.ihelse.net/hl7v3/infrastructure/rim/rim.html#Act-actionNegationInd-att. (Accessed 7/24/2018)

[8] http://hl7.ihelse.net/hl7v3/infrastructure/rim/rim.html#Act-actionNegationInd-att. (Accessed 7/24/2018)

[9] https://confluence.ihtsdotools.org/display/DOCEG/2.4.13.1+Situation+with+Explicit+Context+Attributes+Summary. (Accessed 7/30/2018)

[10] For example, see http://www.hl7.org/fhir/us/core/ValueSet-us-core-substance.html. (Accessed 7/30/2018)

[11] http://solor.io/. (Accessed 7/24/2018)

[12] https://www.hl7.org/documentcenter/public_temp_E7E64027-1C23-BA17-0C1259137A80413F/wg/inm/datatypes-its-xml20050714.htm#comp-CD.qualifier. (Accessed 7/24/2018)