

The logo for SNOMED International, featuring the word "SNOMED" in a large, bold, white sans-serif font above the word "International" in a smaller, white sans-serif font, both set against a solid blue square background.

SNOMED
International

Leading healthcare
terminology, worldwide

Decision Support with SNOMED CT®

31 March 2017

Latest web browsable version: <http://snomed.org/cds>

SNOMED CT Document Library: <http://snomed.org/doc>

1. Decision Support with SNOMED CT	3
1. Introduction	4
1.1. Overview	5
1.2. Functional Areas	7
1.3. Clinical Areas	8
1.4. SNOMED CT Features	10
1.5. Abbreviations	12
2. Logical Architecture	13
2.1. EHR System Architecture	14
2.2. CDS System Architecture	16
3. Knowledge Base	18
3.1. Rules	19
3.1.1. Context in CDS Rules	20
3.1.2. Rule Components and Criteria	24
3.1.3. Rule Examples	25
3.1.4. Standards for CDS Rules	27
3.2. Guidelines	29
3.2.1. Guidelines with SNOMED CT	30
3.2.2. Standards for CDS Guidelines	33
3.3. Substrate	35
4. Inference Engine	36
4.1. Reasoning with SNOMED CT	37
4.1.1. Reasoning with Subsets	38
4.1.2. Reasoning using Subsumption	40
4.1.3. Reasoning using Defining Relationships	42
4.1.4. Reasoning with Description Logic	44
4.2. Accessing Clinical Records	45
4.2.1. Standards for Accessing Clinical Records	46
4.3. Accessing Terminology	48
5. Communications	50
5.1. Standards for CDS Communications	52

Decision Support with SNOMED CT

The Decision Support with SNOMED CT guide reviews the approaches, tools and techniques used to implement clinical decision support with SNOMED CT, and shares developing practice in this area. It is anticipated that this guide will benefit members, vendors and users of SNOMED CT by promoting a greater awareness of how SNOMED CT has been and can be used to enhance clinical decision support implementations.

Latest web browsable version: <http://snomed.org/cds>

Directory of available documents: <http://snomed.org/doc>

Publication date: 2017-03-30

This is the date when this PDF file was generated from the online version of this document.

© Copyright 2017 International Health Terminology Standards Development Organisation, all rights reserved.

This document is a publication of International Health Terminology Standards Development Organisation, trading as SNOMED International. SNOMED International owns and maintains SNOMED CT®.

Any modification of this document (including without limitation the removal or modification of this notice) is prohibited without the express written permission of SNOMED International. This document may be subject to updates. Always use the latest version of this document published by SNOMED International. This can be viewed online and downloaded by following the links on the front page or cover of this document.

SNOMED®, SNOMED CT® and IHTSDO® are registered trademarks of International Health Terminology Standards Development Organisation. SNOMED CT® licensing information is available at <http://snomed.org/licensing>. For more information about SNOMED International and SNOMED International Membership, please refer to <http://www.snomed.org> or contact us at info@snomed.org.

1. Introduction

Background

SNOMED CT is a standardized and multilingual clinical terminology used by clinicians and other health care providers to record and share health information. As the most comprehensive terminology in the world, SNOMED CT contains over 300,000 active clinical concepts, each representing a unique clinical meaning. These concepts are organized into hierarchies such as [Clinical finding](#), [Procedure](#), [Body structure](#) and [Pharmaceutical / biologic product](#).

SNOMED CT is increasingly being used in clinical decision support (CDS) systems to support healthcare providers in making well informed clinical decisions. SNOMED CT's polyhierarchy, defining relationships and concept model are just some of the terminology's features that help to link patient records to the appropriate guidance, clinical knowledge and decision rules.

Purpose

The guide [Data Analytics with SNOMED CT](#) explained how SNOMED CT may be used to support data analytics, and how integrating clinical records with decision support tools can improve the care provided to individual patients by guiding safe, appropriate and effective patient care.

The purpose of this guide is to review the approaches, tools and techniques used to implement clinical decision support with SNOMED CT, and to share developing practice in this area. It is anticipated that this guide will benefit [Members](#), vendors and users of SNOMED CT by promoting a greater awareness of how SNOMED CT has been and can be used to enhance clinical decision support implementations.

Scope

This guide introduces the key components of clinical decision support systems, explores ways in which SNOMED CT can be used to enhance the capabilities within each of these components, and presents some case studies in which SNOMED CT has been used to support clinical decision support. The guide focuses on CDS systems that use a combination of SNOMED CT encoded knowledge artifacts (e.g. CDS rules and guidelines) and SNOMED CT encoded electronic health records. However, using SNOMED CT to enhance non-SNOMED CT enabled CDS and EHR systems is also considered.

Audience

The target audience of this guide includes:

- Members who wish to learn about using SNOMED CT for clinical decision support
- Clinicians, informatics specialists and technical staff involved in the planning, management, design or implementation of clinical record applications or clinical decision support systems
- Software vendors, data analysts, epidemiologists and others designing SNOMED CT based solutions

This guide assumes a basic level of understanding of SNOMED CT. For background information it is recommended that the reader refers to the [SNOMED CT Starter Guide](#).

Guide Overview

The guide presents an introduction to clinical decision support using SNOMED CT and is structured as follows:

1. **Introduction:** Chapter 1 provides an introduction to the guide, and defines clinical decision support (CDS) and clinical decision support systems (CDSS). It then presents an overview of CDS (including its scope, history and the 'five rights'), it explores the functional and clinical areas in which CDS is used, the features of SNOMED CT that support CDS, and a table of abbreviations used in this guide.
2. **Logical Architecture:** Chapter 2 presents an overview of the logical architecture of an electronic health records system that uses CDS, and the internal components of a CDSS system.
3. **Knowledge Base:** Chapter 3 describes the knowledge base of a CDSS, in which the knowledge artifacts that drive the CDS are stored, and how SNOMED CT may be used within these artifacts.
4. **Inference Engine:** Chapter 4 explains how the inference engine of a CDSS can use SNOMED CT to execute the knowledge artifacts in the knowledge base.
5. **Communications:** Chapter 5 explores some of the considerations around implementing the communication mechanisms in a CDSS.

1.1. Overview

What is Clinical Decision Support?

Clinical Decision Support (CDS) is a service that enables healthcare providers to make well-informed decisions by supplying guidance, knowledge, and patient-specific information at relevant points in the patient journey, such as diagnosis, treatment, and follow-up. CDS uses a range of mechanisms to assist users in this process. Examples of these mechanisms include automated alerts or reminders, clinical guidelines, contextually relevant reference information, conditional order sets, diagnostic support, and patient-focused reports, forms, or templates. The beneficiaries of the information derived from CDS may include patients, clinicians, and others involved in the delivery of health care.

It is important to distinguish the general practice of clinical decision support from the application of tools designed to enhance decision support practices. One is performed by humans who make decisions based on knowledge they possess and information they consume. The other is computed by systems and engines using rules and predefined conditions. Although both are important, the technical components of CDS are designed to assist rather than replace the subtle judgment and guidance provided by the clinician.

Applications and tools that provide clinical decision support are known as Clinical Decision Support Systems^[1] (CDSS) and these are defined as follows:

A **clinical decision support system** is defined as a computer system or software application designed to support clinicians, other health professionals, carers or patients making decisions related to the health and treatment of a patient.

Notes

1. Typically a clinical decision support system responds to triggers, such as specific symptoms, signs, diagnoses, laboratory results, medication choices, or complex combinations of these. The system then provides information or recommendations directly relevant to the specific patient.
2. **Clinical decision support** (CDS) refers to services provided (or potentially provided) by clinical decision support systems.

History

It has been suggested that the origins of clinical decision support (CDS) can be traced back to the 1950s and 1960s. In an early example from 1961, Dr. Homer Warner, a cardiologist from the University of Utah developed a mathematical model which was used to diagnose heart disease.^[2] Since then, there have been countless developments and advancements in the area of decision support. Many theories have been proposed as to how CDS should be approached and applied in clinical practice.

The Five Rights

When implemented properly, CDS has the potential to enhance patient care, reduce errors and duplication of effort, and introduce efficiencies to the clinical workflow. Conversely, CDS tools can also be distracting and disruptive, even producing unwanted consequences. It is therefore important to consider the lessons learned from previous implementations of CDS and conduct thorough requirements analysis prior to designing or procuring a CDSS. One of the best practice frameworks that has been developed to guide those considering a CDS implementation is the "CDS Five Rights".^[3] "The Five Rights" suggests that to realize the full potential of CDS, solutions should:

- Supply the right information (evidence-based guidance, address the clinical need)
- To the right people (entire care team, including the patient)
- Using the right channels (e.g., EHR, mobile devices, patient portals)
- In the right intervention formats (e.g., order sets, flow-sheets, dashboards, patient lists)
- At the right points in the workflow (for decision making or action)

Example

A typical application of CDS is shown in the diagram below:




Figure 1.1-1: Example of simple application of CDS

The clinical setting in which this hypothetical tool has been applied is the prescribing of a medication. In this example, the patient has previously had an [Allergy to penicillin](#) recorded. When prescribing a new drug, such as [Amoxicillin](#), an alert is displayed to remind the clinician of the previously diagnosed allergy. The application may also provide a mechanism to search for alternative medications. Note that the mechanics of this workflow uses a predefined rule which specifies a condition to be evaluated and an action to be taken if the condition evaluates to true.






- 1 Definition based on content from U.S. Department of Health and Human Services, Agency for Healthcare Research and Quality
- 2 Timeline of the Development of Clinical Decision Support (OHSU)
- 3 Osheroff, Teich, Levick et al., 2012. Improving outcomes with CDS: an implementer's guide, Second Edition

1.2. Functional Areas

This section addresses the functional scope of clinical decision support. CDS may be represented in a variety of formats or tools which depend on the clinical situation or environment. These tools are often referred to as CDS formats, types, or interventions and can be deployed to a wide variety of systems and platforms, such as mobile devices.

Within these functional areas, CDS can be further subdivided into tools which are prompted by the tasks a clinician performs such as patient charting or diagnosis, and functions that are triggered by external events such as the expiration of a period of time.¹ Some of the more common CDS functions are described briefly in the table below. Use of these functions may be appropriate in a variety of clinical domains or use cases, some of which are discussed in the [Clinical Areas](#) section.

Table 1.2-1: CDS Functions

CDS Function		Description
	Alerts or Reminders	One of the more common types of decision support is computerized alerts (or reminders). These are triggered by rules and designed to interrupt clinicians or patients at the appropriate time. These alerts are also referred to as "best practice advisories" and can be implemented as pop-ups on a users screen or in monitoring tools such as a dashboard. Alerts can also be used to trigger other communication mechanisms such as paging or faxing. Examples of alerts include drug to drug interactions, or drug allergy warnings triggered when medications are prescribed.
	Clinical Guidelines and Reference Information	These CDS functions are often implemented as links to external references which are published by third party, knowledge experts. Guidelines may be represented in a standardized format to facilitate interoperability - for example, the HL7 Infobutton ² . References can be based on relevant, context-dependent data captured in a patient health record or another electronic artifact such as an order or clinical document.
	Diagnostic Support Tools	These tools use a combination of patient data, context-based suggestions and clinical knowledge links to aid the clinician in making a diagnosis. An example would be a tool that prompts a physician for additional findings and suggests additional tests or procedures to help differentiate the diagnosis.
	Automatically Triggered Smart Forms	These documentation tools, which include reports and summaries, are aimed at high quality records, the reduction of errors, and more complete information. These tools can be triggered when a specific patient condition is detected or when a finding is deemed reportable to a jurisdictional health body. These can be represented as focused patient data reports or summaries and are often utilized at the point of care (POC) in real time.
	Conditional Order Sets and Pathway Support	These are typically designed for complex ordering scenarios. They may be comprised of a proposed set of orders or a treatment regimen which is based on an explicit situation or medical condition. These interventions can ensure compliance with established protocols. They can also be utilized to guide clinicians through complex care pathways.





¹ <http://www.himss.org/library/clinical-decision-support/issues>

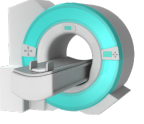





² http://www.hl7.org/implement/standards/product_brief.cfm?product_id=208

1.3. Clinical Areas

The focus of this section is the clinical application of CDS tools or how the [functional components](#) described earlier can be used in practice. Stakeholders from various clinical domains interact with clinical systems, such as EHRs with CDSS and CPOE (computerized physician order entry). The table below lists some of the clinical areas in which SNOMED CT enabled CDSSs can assist clinicians in making well informed decisions.

Table 1.3-1: Clinical Areas

Clinical Area		Description
	Medication Management	A clinician uses an EHR with CDS to prescribe Warfarin sodium 4mg tablet . The CDSS queries the EHR and discovers that the patient is Pregnant . The CDSS determines that the proposed drug has Warfarin as an ingredient. As warfarin is contraindicated during pregnancy, the system triggers an alert to be displayed to the clinician. Relevant clinical guidelines are also displayed to the user. These guidelines suggest a safe alternate, such as Dabigatran , which the clinician then safely prescribes to the patient.
	Diagnosis (e.g. Diabetes) ¹	A clinician uses an EHR with CDS in a case analysis scenario to aid in diagnosis. The clinician records the patient's age and gender, then prepares to enter specific clinical findings, history, symptoms, etc. As the physician records symptoms of Hunger , Fatigue , and Dry mouth , a ranked list of common diseases, associated with these clinical findings, is dynamically presented to the clinician. At the top of this list is Diabetes mellitus . A scale is used to indicate the level of support for each disease. The CDSS then prompts the clinician for additional findings to help differentiate between diseases. Once a confirmed diagnosis is made, the differential diagnoses can be marked as Absent , Present , or Unknown . An additional finding of Always thirsty is recorded and the level of support for each disease in the list is adjusted accordingly. Support for Diabetes mellitus has now increased from minimal evidence to sufficient evidence. The clinician then selects Type 2 diabetes mellitus which opens an evidence screen displaying the recorded findings which either strongly support, support, or do not support the chosen disease. The clinician is then presented with a link that displays all the PubMed articles associated with Type 2 diabetes mellitus .
	Laboratory (e.g. Critical Results) ²	A patient presented at Emergency complaining of Chest pain and was subsequently admitted to the hospital. The attending physician ordered a series of lab tests including a Serum potassium measurement . Laboratory tests are completed and published to the laboratory information system (LIS). The CDSS then queries the LIS and learns that the Potassium level is Low serum potassium level and considered critical. The CDSS then queries the EHR to confirm the patient has been prescribed Oral form digoxin , which has Digoxin as an active ingredient. A knowledge base rule has been defined which stipulates, if the drug prescribed contains Digoxin and the laboratory test indicates a Low serum potassium level , then inform the user. An alert, in the form of an urgent pager message, is generated and sent to the attending physician.
	Radiology (e.g. Contraindication)	An ordering physician has requested an upper Gastrointestinal tract x-ray , which uses Barium sulfate materials. The patient presents at the imaging clinic on the day of their exam. During study protocoling, the imaging department uses the CDSS to query the patient record and determine the patient has a History of hay fever . An alert is triggered to advise the imaging technician about the risk of an allergic reaction. The imaging department, in consultation with the GI radiologist, calls the ordering doctor to discuss the associated risks. Additional guidelines related to preparing for reactions and symptom management (Hives , Itching , Swelling , etc.) are provided via the CDSS. An additional medication is administered prior to the contrast material to reduce the risk of an allergic reaction. The imaging department proceeds with the planned procedure.

	<p>Radiology</p> <p>(e.g. Appropriate Imaging)</p>	<p>A clinician records notes into the appropriate fields of an EHR. For example, Clinical notes: "Pt is 75 yo. LBP (lower back pain) for the past 2 weeks. On exam normal SLR (straight leg raise)..." Using NLP, these notes are encoded as part of the record storage process. (For example, as Low back pain and On examination - straight leg raising normal - left right .) The clinician orders a series of imaging tests. The CDSS, based on specific quality metrics (e.g., appropriate use criteria or AUC), evaluates whether or not imaging guidelines are being followed by analyzing the patient's health record together with the proposed tests. If the guidelines were not followed, the CDSS will display an alert informing the clinician that they may want to consider alternative imaging or additional tests. For example, an alert may indicate: "The patient has Low back pain and Numbness of lower limb . A MRI of lumbar spine without contrast for this case has an appropriateness rating of 8 (scale of 10) and is recommended."</p>
	<p>Emergency Department</p> <p>(e.g. Order sets)</p>	<p>A patient has presented at the Emergency Room (ER) complaining of Shortness of breath . The attending physician records the appropriate clinical finding codes in the EHR. She then prepares a condition-specific order set in a Computerized Physician Order Entry (CPOE) system. The selection of the order set triggers the presentation of new clinical guidelines based on an analysis of the patient record with the proposed treatment. The physician then chooses alternative treatment. Suggested dosage guidance is provided by relevant contextual links within the order set.</p>
	<p>Infectious Disease Reporting</p>	<p>A primary care physician logs on to their EHR with CDS and opens a patient chart to record a condition deemed communicable, such as Mumps or Measles . The CDSS then triggers an alert to advise the provider that this condition is considered reportable to the jurisdictional public health office. The CDSS then provides a pre-populated smart form which facilitates quick, consistent, and accurate reporting of the condition to the local officer of medical health. The smart form is completed and submitted to the jurisdictional health office. The clinical findings in the report are terminology-encoded which promotes interoperability and facilitates population based health reporting.</p>
	<p>Clinical Treatment Audit</p>	<p>A department head uses an EHR with CDS to conduct a treatment analysis. She uses the system to generate a list of all inpatients with a confirmed diagnosis of Deep venous thrombosis . She then uses the system to determine which of these patients have received Heparin therapy for at least 72 hours. The patients which have not met this criteria are flagged for appropriate treatment.</p>
	<p>Acute Asthma Management</p>	<p>Staff in an Emergency Department (ED) use their EHR with CDS and clinical management pathways to provide a standardized evidence-based approach to patient assessment of Acute asthma in adults^[3]. The guidelines help document indications and contraindications to determine eligibility. A triage nurse queries the EHR and learns that the patient is over 16 years of age, has a Acute asthma , and one or more episode of Wheezing which necessitated Breathing treatment . The CDSS then triggers an alert to follow the pathway's medical directives, which are carried out by a Respiratory Therapist (RT). The directives, in this case of Mild asthma , include Heart rate monitoring , establishing various baseline Respiratory measurements , and administration of a Bronchodilator and Prednisone 50mg tablet . The RT then notifies the attending physician who fills out and signs discharge instructions which a nurse then reviews with the patient. The desired clinical outcomes of this pathway include improved adherence to evidence-based management and improved patient outcomes such as reduced number of hospitalizations and lower ED return rates.</p>
	<p>Nursing Interventions</p>	<p>Research has provided evidence to show that patients receiving Mechanical ventilation are at high risk for Pneumonia : due to = Aspiration . ^[4] Published guidelines recommend Elevation of head of bed from 30° to 45°, if not contraindicated, to reduce risk of Pneumonia . A nursing supervisor uses a dashboard-like tool in an ICU to monitor patients in her ward. Patients who meet the criteria for risk of Aspiration pneumonia are automatically flagged in the system using CDS logic so that the appropriate action may be initiated by nursing staff in the ward. Once the angle of the patient's bed is adjusted, the system is dynamically updated and the flag is removed.</p>

^[1] Based on workflow described in DXPlain product literature.

^[2] Example uses material from Brigham and Women's Hospital, and was cited in the [BC Medical Journal](#).

^[3] Uses material from [Educational and decision-support tools for asthma-management guideline implementation and Adult Emergency Department Asthma Care Pathway](#).

^[4] <http://ccn.aacnjournals.org/content/32/3/71.full>



1.4. SNOMED CT Features

Overview

This section contains a brief summary of key SNOMED CT features and explains how they may be useful in CDSSs.

Concepts

SNOMED CT [concepts](#) are used to represent clinical meanings. Every concept in SNOMED CT is uniquely identified by a distinct SNOMED CT Concept Identifier. For example, 195967001 is the concept identifier for the concept [Asthma](#).

SNOMED CT concepts play an important role in CDS by enabling actions to be triggered based on the meaning of data recorded in the patient records.

Descriptions

SNOMED CT [descriptions](#) provide the human-readable terms associated with SNOMED CT concepts. A concept may have one or more descriptions, which act as synonyms for the same clinical meaning. This is also how SNOMED CT supports different dialects and languages.

SNOMED CT descriptions allow common CDS rules to be consistently applied across patient records recorded using different synonyms, dialects and languages.

Relationships

SNOMED CT [relationships](#) link concepts together to formally define the meaning of each concept. For example, one type of relationship is the [is a](#) relationship which relates a concept to a parent or supertype. These [is a](#) relationships define the subtype hierarchy of SNOMED CT concepts.

For example, the concepts [Bacterial pneumonia](#) and [Viral pneumonia](#) both have an [is a](#) relationship to [Infective pneumonia](#) which has an [is a](#) relationship to the more general concept [Pneumonia](#). Subtype relationships can be used by CDS rules to refer to codes in an EHR that are any specific type of a relevant clinical concept.

Additional attribute relationships help to define the meaning of a concept. For example, the concept [Viral pneumonia](#) has a [Causative agent](#) relationship to the concept [Virus](#) and a [Finding site](#) relationship to the concept [Lung structure](#).

Attribute relationships can be used by CDS rules to refer to codes recorded in an EHR that have a specific meaningful relationship with a concept of interest.

Concept Model

The SNOMED CT [concept model](#) is a set of rules that govern the ways in which SNOMED CT concepts are permitted to be modeled using relationships to other concepts. It defines the types of relationships that may be used on each type of concepts, and the permitted values for each relationship type. The [Machine Readable Concept Model \(MRCM\)](#) represents the rules in the SNOMED CT concept model in a form that can be read by a computer and applied to test that concept definitions and expressions comply with these rules.

The SNOMED CT concept model plays an important role in CDS by providing the rules by which the clinical meaning of SNOMED CT encoded health records can be queried. The MRCM makes it possible to process these rules in a machine-processable way.

Expressions

SNOMED CT provides a mechanism which enables clinical phrases to be represented by a computable [expression](#), when a single concept does not capture the necessary level of detail. For example, the following expression represents a right hip:

```
182201002 |Hip joint| :
  272741003 |Laterality| = 24028007 |Right|
```

SNOMED CT expressions enable additional clinical meanings to be captured in a health record, without requiring the terminology to include countless combinations and permutations of precoordinated concepts.

SNOMED CT expressions facilitate CDS over an expanded set of clinical meanings that extends beyond individual concepts. For more information about expressions, please refer to the [Compositional Grammar - Specification and Guide](#).

Reference Sets

SNOMED CT [reference sets](#) are a flexible and standardized approach used to support a variety of requirements for the customization and enhancement of SNOMED CT. These include the representation of [subsets](#), language preferences for use of particular terms, mapping from or to other code systems, and ordered lists.



Reference sets may be used in the following aspects of CDS:

- Representing subsets of SNOMED CT concepts that may trigger a CDS action
- Representing non-standard aggregations of concepts for specific CDS use cases
- Defining language or dialect specific sets of descriptions over which term searches can be performed

For more information about reference sets, please refer to the [Practical Guide to Reference Sets](#).

Description Logic Features

Description Logic (DL) is a family of formal knowledge representation languages and used as the formal foundation of meaning in SNOMED CT. The way that concepts have been modeled in SNOMED CT permits them to be represented using Description Logic. DL helps computers to make useful inferences about concepts, and to classify SNOMED CT using a DL reasoner. Description Logic also helps by testing expressions for subsumption and equivalence.

The logical inferences supported by DL can be useful when executing CDS rules. For example, when a CDS rule requires an action to be performed when the patient has any type of `|Asthma|`, a DL reasoner may be used to determine that `|Acute asthma|` and `|Intermittent asthma|` are both types of `|Asthma|` and should therefore both trigger the action to be performed.

1.5. Abbreviations

The following table contains the definition of abbreviations used in this document. Please refer to the [SNOMED Glossary](#) for additional definitions.

Table 1.5-1: Abbreviations

Abbreviation	Full term linked to the SNOMED Glossary definition
CDS	Clinical decision support , which is defined as a service that supports clinicians, other health professionals, carers or patients making decisions related to the health and treatment of a patient.
CDSS	Clinical decision support system , which is defined as a computer system or software application designed to support clinicians, other health professionals, carers or patients making decisions related to the health and treatment of a patient.
EHR	Electronic health record , which is defined as a systematic collection of health information about individual patients or populations that is stored in a digital form. An Electronic health record may contain a complete and detailed record of a patient's health or may consist of a summary of information of particular relevance to continuing delivery of care.
KB	Knowledge base , which is defined as the underlying set of facts, assumptions, and rules which a computer system has available to answer a question or solve a problem.
UI	User interface , which is defined as the way a software application presents itself to a user including, its on screen appearance, the commands it puts at a users disposal, and the manner in which the user can access and update information by using the application.
NLP	Natural language processing , which is defined as a service in which a computer system converts between human-readable text (and/or spoken languages) and formal representations of information that can be readily generated, analyzed and processed by other software applications.
POC	Point of care , which is defined as the time and location at which clinicians or other health professionals deliver healthcare products and services to patients.

2. Logical Architecture

This section provides an overview of the logical architecture of an electronic health record (EHR) which uses CDS.

In particular, it focuses on the logical architecture of knowledge-based CDSSs,^[1] which use pre-loaded CDS artifacts (such as rules and guidelines) that closely match a human's natural reasoning process. Non-knowledge-based CDSSs, which use artificial intelligence (AI) or machine learning to acquire knowledge over time, are outside the scope of this guide.

The logical architecture of knowledge-based CDSSs are explored in the following subsections:

- 2.1. EHR System Architecture
- 2.2. CDS System Architecture

[1] https://en.wikipedia.org/wiki/Clinical_decision_support_system#Knowledge-based_CDSS





2.1. EHR System Architecture

To understand how clinical decision support (CDS) works, it is important to understand how CDS fits within the logical architecture of an EHR system.

This section describes the major architectural components of an EHR system, and the interactions between these components.

Major Components

Table 2.1-1: Descriptions of the major architectural components of an EHR system

EHR System Component		Description
	User Interface	The user interface (UI) is a fundamental component of almost any clinical application, and is used within an EHR to both enter and display patient health records. The UI also has two main CDS functions. Firstly, the UI is used to provide inputs to the CDSS, such as recording a proposed medication or an observed finding. The second function of the UI is to display alerts, advisories and clinical guidelines to the user in an appropriate format, on behalf of the CDSS.
	Record Services	Record services are a set of services for managing patient health records. Record services provide functions like entering data into health records, searching for and retrieving health records, querying or extracting data from health records, and communicating or exchanging health records with other systems or applications. Record services interact with other components in this model such as the CDSS and the UI.
	Terminology Services	Terminology services are those services that directly manage the terminology resources. They include functions like querying concepts, relationships and reference sets, and installing or updating SNOMED CT from release files. Terminology services interact with the CDSS component in this model.
	Clinical Decision Support System	The primary role of the CDSS is to execute the decision support logic. The CDSS does this using a number of subcomponents and subprocesses, which will be described in the next section - <i>Logical Architecture of a CDSS</i> . The CDSS interacts with each of the other major components in the EHR.

Interactions

The major architectural components of an EHR system that incorporates CDS interact with each other in a variety of ways to support the overall functioning of the EHR.

The diagram below uses orange arrows to illustrate the primary interactions between these EHR components.

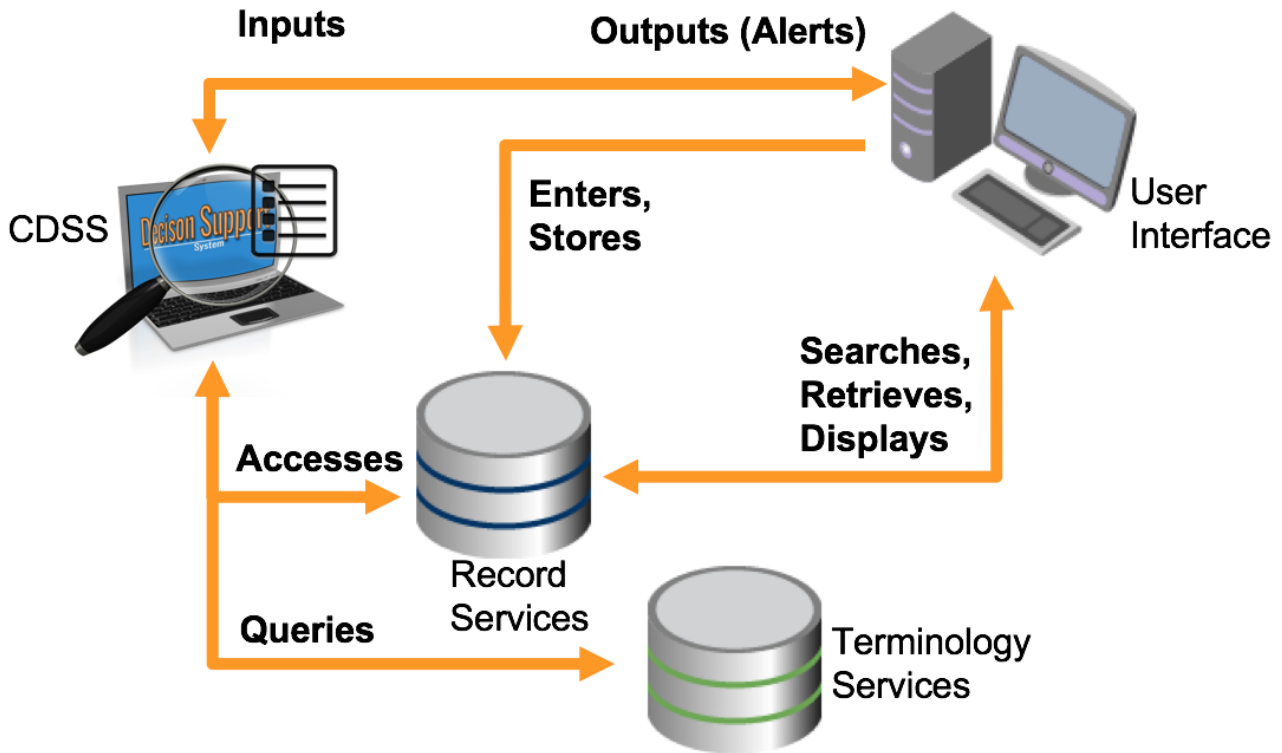


Figure 2.1-1: EHR components and key interactions




As shown above, the UI communicates with the record services to facilitate the storage and subsequent retrieval of health record data. The UI also provides inputs to the CDSS and displays alerts and guidelines on its behalf. The CDSS uses the inputs from the UI and data from the record and terminology services to process decision support rules. The CDSS uses the inputs from the record and terminology services to determine whether or not the CDS conditions have been met, and if so then CDS interventions, such as alerts or knowledge resources, are delivered back to the UI. The internal components and processes of the CDSS will be described in more detail in the next section - Logical Architecture of a CDSS.

2.2. CDS System Architecture

This section describes the major architectural components of a CDSS and explores how they work together with the components of an EHRs, as described in 2.1. EHR System Architecture.

Major Components

Table 2.2-1: Descriptions of the major architectural components of a CDSS

CDSS Component		Description
	Knowledge Base	<p>The knowledge base (KB) stores clinical knowledge developed by domain experts as CDS artifacts. These knowledge artifacts (e.g. rules and guidelines) are stored in a machine processable format and made available to the inference engine to drive the CDS workflow.</p> <p>For additional information on this component please refer to section 3. Knowledge Base.</p>
	Inference Engine	<p>The inference engine processes the CDS knowledge artifacts, using information from the record services, the terminology services and user input to execute the CDS logic. A key part of this process is to determine which actions should be performed, based on the given patient's circumstances.</p> <p>For additional information on this component please refer to section 4. Inference Engine.</p>
	Communications	<p>The communications mechanism is responsible for accepting inputs from the user and delivering the outcomes of the inference engine back to the user. For example, when a clinician prescribes a drug, this information is communicated to the inference engine as an input. If the inference engine discovers that the medication is contraindicated, the communications mechanism will deliver an alert to the user interface.</p> <p>For additional information on this component please refer to section 5. Communications.</p>

Logical Architecture

The diagram below illustrates how the components of the CDSS (shown in the blue box) work together with the components of the EHR system (shown in the red box).

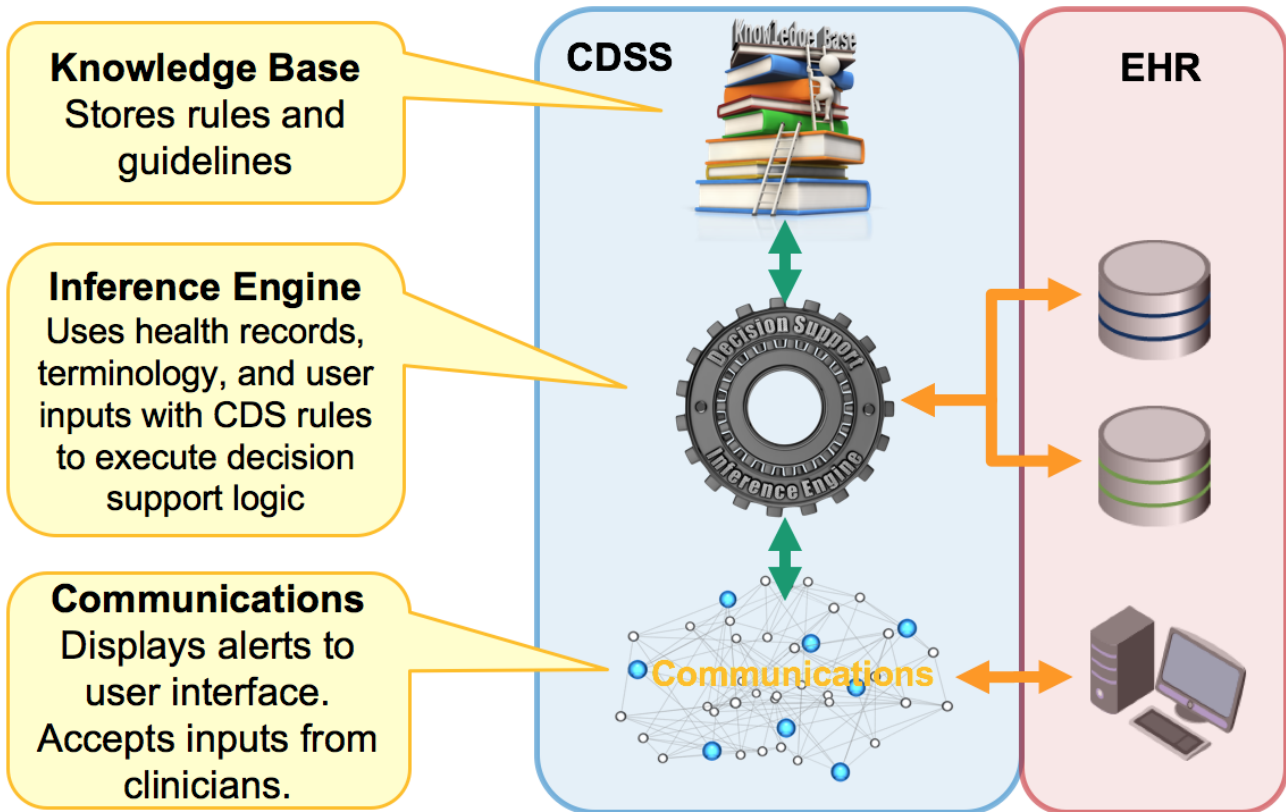


Figure 2.2-1: CDSS components and key interactions

Internal CDSS interfaces are represented by the green directional arrows, while external CDSS interfaces are represented by the orange arrows. Note that the inference engine interfaces directly with record and terminology services while communications, which is focused on the delivery of CDSS inputs and outputs, interfaces directly with the user interface.

3. Knowledge Base

The knowledge base can be thought of as the brains of a clinical decision support system. Clinical knowledge is what fuels the knowledge base. This knowledge is developed by clinical experts in their respective domains. The knowledge is then loaded into the KB as knowledge artifacts and stored in a machine processable format. These artifacts are then made available to the *inference engine* to execute the decision support logic. Knowledge artifacts may be updated when new clinical knowledge becomes available. In some CDSSs this is done using a specialized knowledge artifact management interface, which may support tasks such as rule creation, customization, and updating. In other cases, clinical knowledge artifacts are used from third party providers, who specialize in supporting CDSSs.

Types of CDS knowledge artifacts include:

- Decision support rules
- Clinical guidelines and care pathways
- Documentation templates
- Order sets

The characteristics of these knowledge artifacts are described in the section 1.2. *Functional Areas*.

The diagram below illustrates the key interactions with the knowledge base, as described above.

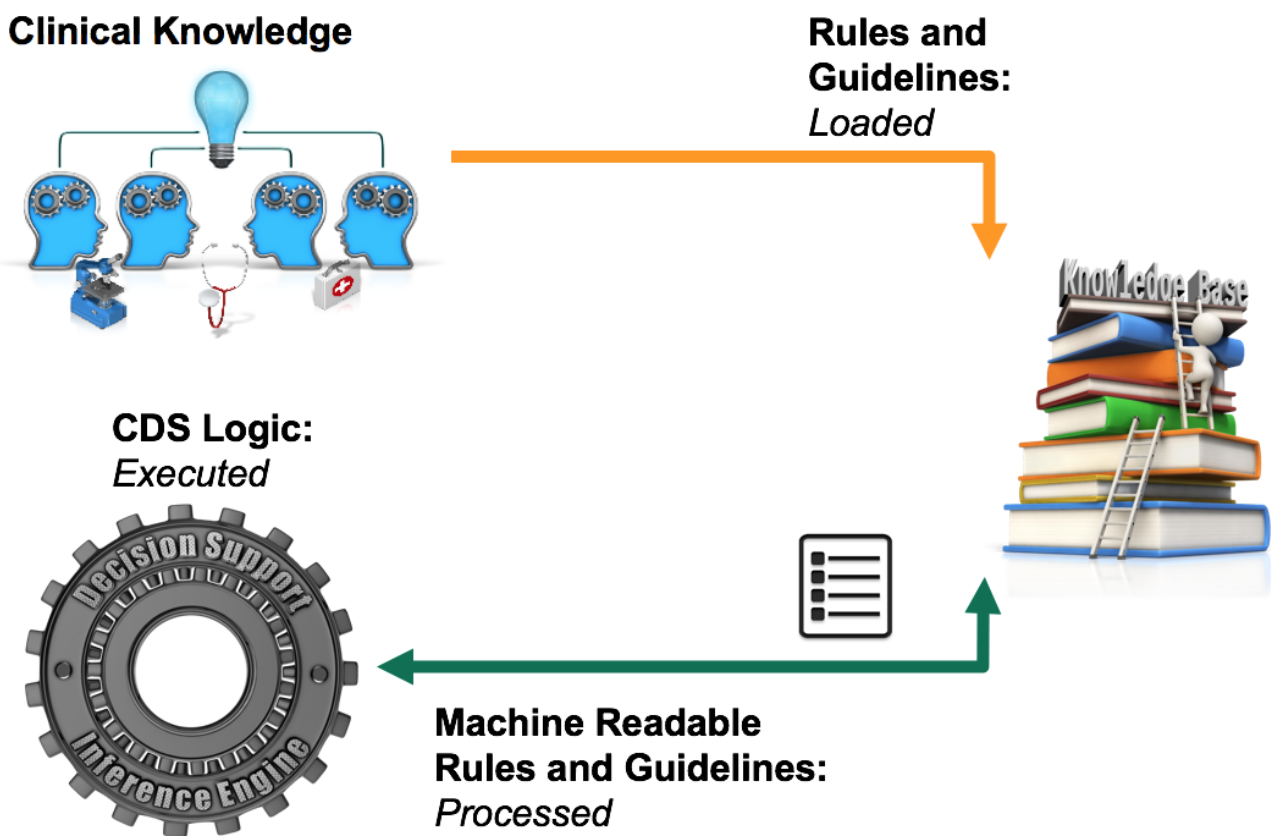


Figure 3-1: Knowledge base interactions

The topics below are presented in more detail in the following sections:

- 3.1. Rules
- 3.2. Guidelines
- 3.3. Substrate

3.1. Rules

Clinical decision support rules play a key role in the overall delivery of CDS. CDS rules typically follow a common pattern, which has been modeled in several healthcare standards formalisms. The Event-Condition-Action¹ model, as used in the HL7 community, is described below.

Event

A CDS event is the clinical situation in which a decision support rule will be applied. First something must happen before the rule can be utilized. Examples of CDS events include:

- A clinician is prescribing a drug to a patient
- A nursing supervisor is reviewing a list of patients previously diagnosed with cancer
- A clinician is assessing a patient enrolled in a jurisdictional diabetes monitoring program

Condition

A CDS condition defines the question(s) that must be answered to determine the outcome of the rule. Examples of conditions include:

- Does the usual drug of choice for this patient's condition contain a substance to which the patient is allergic?
- Have any patients with a suspected cancer diagnosis NOT been referred to a specialist within 14 days of diagnosis?
- Has the patient with a previous diagnosis of diabetes type II NOT had HBA1C tested within the last 12 months?

Action

The CDS action describes what should be done if the condition evaluates to true. Examples of actions include:

- Alert the clinician and suggest a safe alternative medication
- Refer patients to an oncology specialist
- Order HBA1C test

Event-Condition-Action Model

An informal representation or rule template which captures the Event-Condition-Action pattern is shown below. This pattern can be read as "ON event IF condition THEN action".



Figure 3.1-1: Event-Condition-Action rule template

Rules may reference both EHR data and reference data such as terminology to determine whether or not a specific condition is true. This topic will be explored in more detail in section 4. [Inference Engine](#).

¹ <http://hl7.org/fhir/2016Sep/cqif/cqif-knowledge-artifact-representation.html#event-condition-action-rule>

3.1.1. Context in CDS Rules

Context has been defined as the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood. Contexts that modify the meaning of a diagnosis or procedure may include family history, past history, suspected diagnoses, planned procedures and procedures not done. It is important to understand the context of each statement in a health record, to determine whether or not it is appropriate to for a CDS rule to be applied.

When evaluating the condition within a CDS rule it is important to take account of context.

- For example, a rule that requires a current diagnosis of diabetes should not trigger an action in response to a record that states that a patient has a family history of diabetes.

Representing Context in a Health Record

Context can be expressed in a health record in a number of ways. Firstly, a precoordinated expression can be used in which the context is captured in the meaning of the concept. For example, |Family history: Diabetes mellitus|. Alternatively, a postcoordinated expression can be used. This is where the meaning is expressed by combining codes in a structured way using SNOMED CT Compositional Grammar. For example:

```
281666001 |Family history of disorder| :
246090004 |Associated finding| = 73211009 |Diabetes mellitus|
```

A third way to express context is to use a context-specific section or field, such as a "Family history section", which captures the context in the meaning of the section or field name. Lastly, it is also possible to use two separate fields - one which captures the finding |Diabetes mellitus|, and the other which captures the context |Family history of disorder|.

Table 3.1.1-1: Techniques for recording context in an EHR

Technique for Representing Context	Example				
Precoordinated as a single SNOMED CT concept identifier explicitly representing family history of diabetes mellitus.	160303001 Family history: Diabetes mellitus				
Postcoordinated as a SNOMED CT expression that includes a concept representing a family history of disorder and specifies the diabetes mellitus as the disorder.	281666001 Family history of disorder : 246090004 Associated finding = 73211009 Diabetes mellitus				
A context specific family history section in the record structure	<table border="1"> <tr> <td>Family History Record Section</td> <td>73211009 Diabetes mellitus </td> </tr> </table>	Family History Record Section	73211009 Diabetes mellitus		
Family History Record Section	73211009 Diabetes mellitus				
A separate field in the record structure to indicate the context of the disorder recorded	<table border="1"> <thead> <tr> <th>Disorder</th> <th>Context</th> </tr> </thead> <tbody> <tr> <td>73211009 Diabetes mellitus </td> <td>281666001 Family history of disorder </td> </tr> </tbody> </table>	Disorder	Context	73211009 Diabetes mellitus	281666001 Family history of disorder
Disorder	Context				
73211009 Diabetes mellitus	281666001 Family history of disorder				

False Positives and False Negatives

Considering context that is captured in either the terminology or the information structure is important when executing CDS rules.

- Logic based purely on the presence or absence of codes, without considering the context implied by the information structure, may lead to CDS alerts being triggered unnecessarily (i.e. false positives).
- Conversely, logic based purely on the presence or absence of codes, without considering context implied by the information structure, may lead to CDS alerts not being triggered when required (ie. false negatives).

False Positive Example

In the following example, a CDS rule is triggered inappropriately (i.e. false positive):

- A CDS rule is designed to display clinical practice guidelines for stage 1 chronic kidney disease when the code 431855005 |Chronic kidney disease stage 1| is found in the EHR. A retrospective analysis of a false positive trigger reveals that the code was recorded in the past history section of the health record. As this record indicates that the |Chronic kidney disease stage 1| was part of the patient's |Past medical history|, the display of stage 1 chronic kidney disease guidelines was inappropriate.

False Negative Example

In the following example, a CDS rule is not triggered when required (i.e. false negative):

- A CDS rule is designed to display patient-focused, preventative educational material when the code 160303001 |Family history: Diabetes mellitus| is found in the EHR. This rule is implemented in an EHR system, which uses a family history section to record the family history of disorders. Even though the SNOMED CT concept 73211009 |Diabetes mellitus| is recorded in the patient's family history, the CDS rule is not triggered as required.

Default Context

When neither the SNOMED CT concept nor the surrounding health record explicitly states the context, a default context applies.

Table 3.1.1-2: Default context values alongside their corresponding attributes

	Attribute	Value
Clinical Findings	Finding context	Known present
	Subject relationship context	Subject of record
	Temporal context	Current or specified time
Procedures	Procedure context	Done
	Subject relationship context	Subject of record
	Temporal context	Current or specified time

For clinical findings, it is assumed that the finding is known to be present (as opposed to known to be absent), we assume that the finding is about the patient (as opposed to someone else), and we assume that the finding occurred at either the present time or a time specified in the record structure (as opposed to a general time in the past).

Data Entry with Context

The following diagram illustrates how additional context can be captured during the data entry process. The diagnosis (disorder) of |Female breast cancer| is selected using the pick-list in the top left corner of the screen and then the context values are selected using the other radio button and pick-list. These context values for relation to subject and finding context must be considered when the conditions in this CDS rule are evaluated.




Figure 3.1.1-1: Capturing context during data entry

Defining Context in Rules

It is important to always consider context when defining (and executing) the conditions in a CDS rule. If the default context applies to the condition, then it does not need to be explicitly stated in the CDS rule. However, care should be taken when testing the CDS condition against health records, to ensure that the recorded values share the same context as is required by the CDS rule. If a non-default context is required in a CDS rule, then the rule must explicitly state the context that is required. This context must also be appropriately checked when testing the CDS condition against health records.

The following diagram illustrates a CDS rule, which explicitly states the context of a `Female breast cancer` diagnosis that must be matched in order for the action to be triggered. In this example, the CDS condition requires that for patients over the age of 30 `years`, a diagnosis of `Female breast cancer` must be present, in a female family member of the subject, with genetic ties.

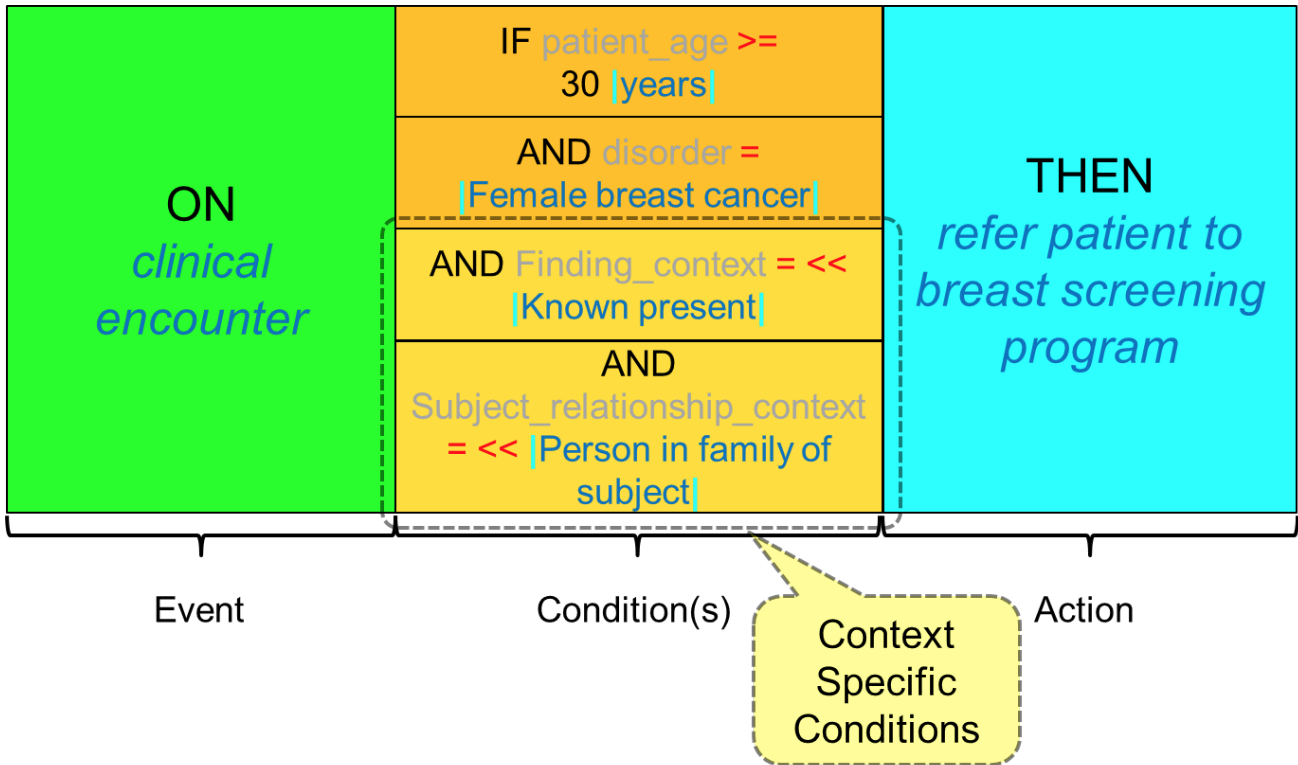


Figure 3.1.1-2: CDS rule with the context explicitly stated

The contextual selections in the data entry screen above would satisfy the conditions in this rule because the user has specified that the diagnosis of female breast cancer occurred in the mother of the subject. This can be seen in the postcoordinated expression below, which corresponds to the user's selections in the data entry screen:

```
372064008 |Female breast cancer| :
  408729009 |Finding context| = 410515003 |Known present| ,
  408732007 |Subject relationship context| = 444301002 |Mother of subject|
```

Note that a selection of |Stepmother| would not trigger the rule as this concept is not a descendant of |Person in family of subject|. (Stepmother has no genetic relationship to the patient.)

<https://en.oxforddictionaries.com/definition/context>

3.1.2. Rule Components and Criteria

The components and criteria within a CDS rule should also be considered when designing or implementing the rules. Some of the additional aspects of these considerations have been described below.

Multi-Component CDS Rules

Multiple events, conditions, or actions may be associated with each CDS rule. For example, two separate actions defined in a medication allergy rule might be:

1. Firstly to alert the user and;
2. Secondly to suggest an alternative drug

Additional examples of complex rules, with multiple conditions and actions, are provided in the section [3.1.3. Rule Examples](#).

Criteria in CDS Conditions

Each CDS condition can be further subdivided into one or more criterion, each consisting of a "name-value" pair. The criterion name will typically map to a data element in the electronic health record, while the criterion value is compared with the data that populates this element in the patient's health record.

Table 3.1.2-1: Examples of criterion name-value pairs

Criterion Name	Criterion Value
Pregnancy status	Pregnant
Drug prescribed	Codeine
Hematocrit result	41 %

Criteria Values

Some criteria may refer to the value of coded data elements, while others may refer to the value of non-coded data elements. When a criterion refers to a SNOMED CT encoded data element, the value may be a [SNOMED CT Expression Constraint](#) that defines the permitted subset of concepts that will satisfy this criteria.

Table 3.1.2-2: Examples of criteria which refer to coded data elements

Criterion Name	Criterion Value
Condition	<< 195967001 Asthma
Diagnosis	<< 29857009 Chest pain
Procedure	< 71388002 Procedure : << 363704007 Procedure site = << 20139000 Structure of respiratory system

Criteria that refer to non-coded data elements may use operators that are valid for the given element's data type. For example, criterion that refer to numeric data elements may use standard mathematical operators to restrict the required value.

Table 3.1.2-3: Examples of non-coded data elements

Criterion Name	Data Type	Operator	Value	Units
Wait time	Quantity	>	90	days
Lab result	Quantity	<=	7.5	mmol/L
Birth date	Date	>=	1990/01/01	n/a

3.1.3. Rule Examples

In this section we present a number of examples of CDS rules, using the 'ON event IF condition THEN action' pattern described in section 3.1. Rules.

Asthma Diagnosis

This simple CDS rule was designed to be used during a clinical encounter. If the patient is diagnosed with asthma, the appropriate management guidelines are automatically displayed on the clinician's workstation.




Figure 3.1.3-1: This rule demonstrates the use of a single event, criterion, and action, an expression constraint in the criterion value, and use of the default context.

Medication Order

This rule has been designed to be used when ordering a medication. If the patient is |Pregnant| and the drug has an active ingredient of |Warfarin|, the clinician will be alerted and the CDSS will suggest an alternative blood thinner which does not pose a risk for expectant mothers.

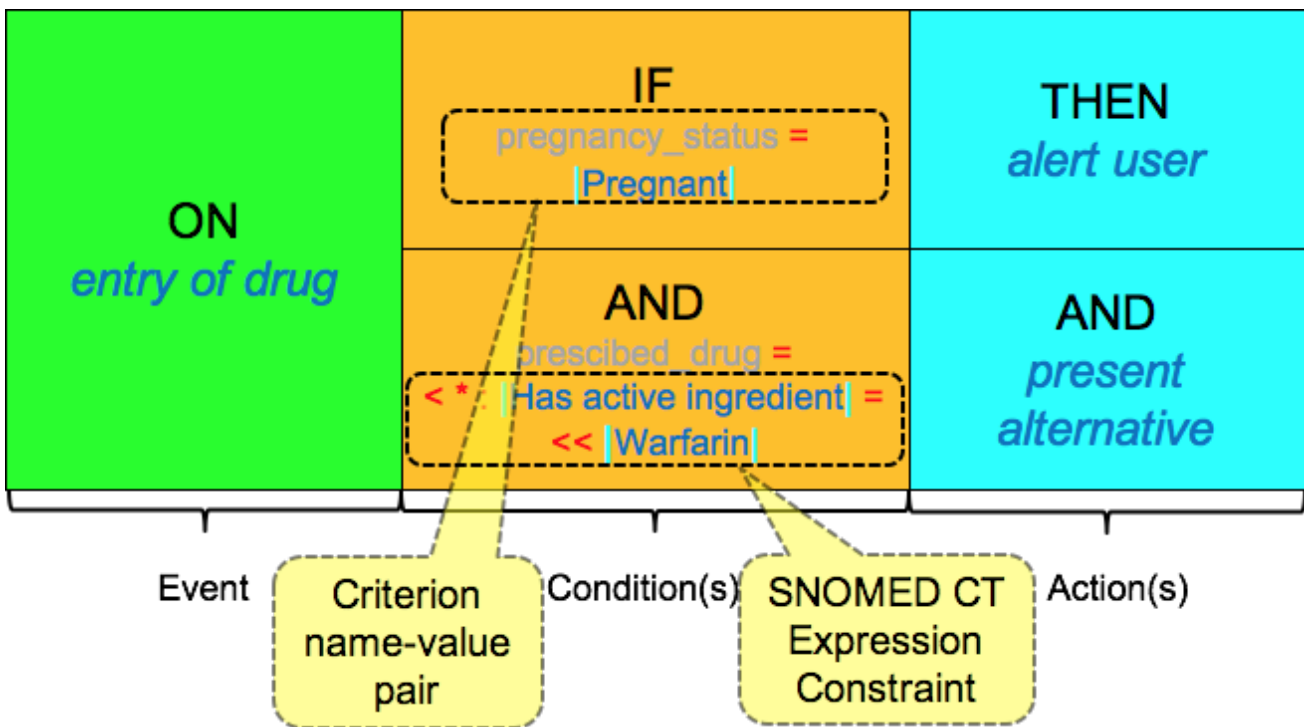


Figure 3.1.3-2: This rule demonstrates the use of multiple criteria, multiple actions, expression constraints in the criterion value, and use of the default context.

Emergency Department

The following example is a CDS rule designed to be used in an Emergency Department setting, when a patient has presented in the ER with chest pain. In this scenario, the attending physician may order a `|Blood potassium measurement|`. If the patient is currently taking a medication with an active ingredient of `|Digoxin|`, and the lab result is published indicating that the patient's potassium level is less than 3.0 mmol/L, the attending physician will be paged.

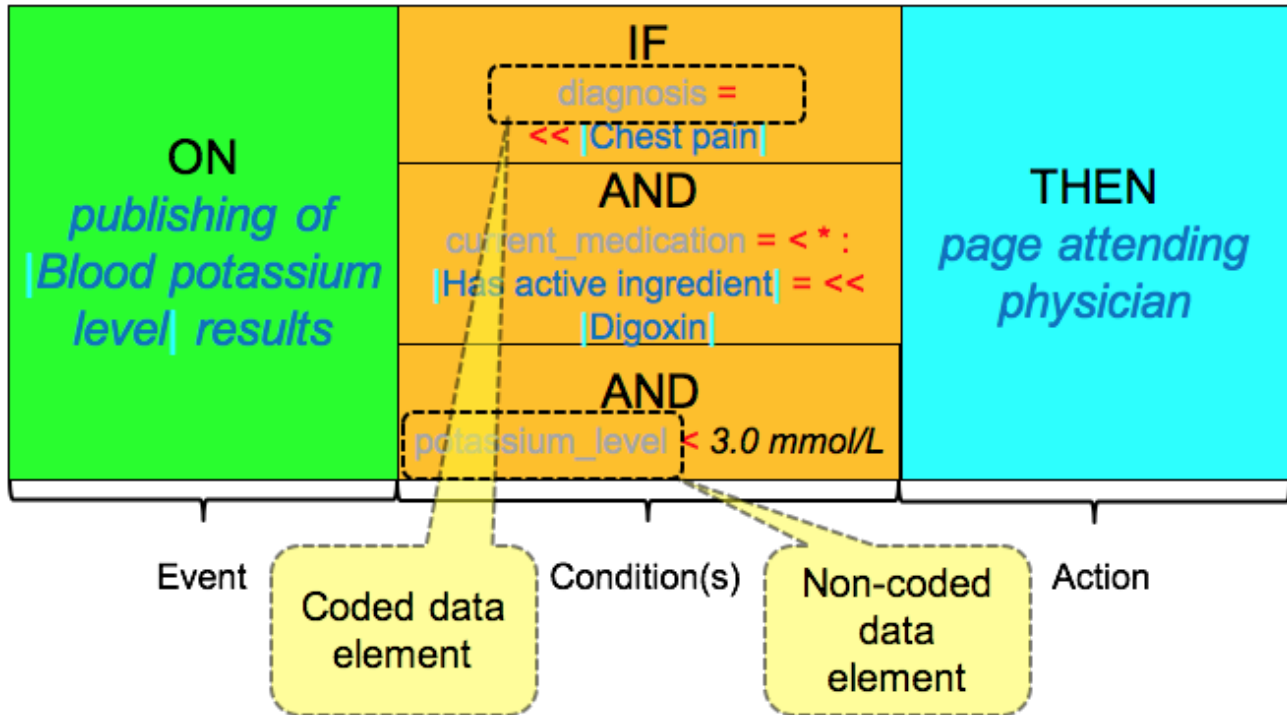


Figure 3.1.3-3: This rule demonstrates the use of multiple criteria, expression constraints in the criterion value, a non-coded criterion with a mathematical operator, and use of the default context.

3.1.4. Standards for CDS Rules

This section presents some examples of standards used to represent CDS rules. Please note that this list is not exhaustive, and other established and emerging standards for rule representation do exist.

Expression Constraint Language

The SNOMED CT [expression constraint language](#) (ECL) provides a computable way of intensionally defining a set of clinical meanings represented in SNOMED CT. For example, the expression constraint below represents the set of lung disorders that have an associated morphology that is a type of edema.

```
< 19829001 |Disorder of lung| :
  116676008 |Associated morphology| = << 79654002 |Edema|
```

When executed against a specific SNOMED CT edition, an expression constraint will return the set of concepts that match the given constraint. Expression constraints can also be used to query over precoordinated and postcoordinated expressions recorded in EHRs.

SNOMED CT expression constraints provide a standard way of referring to intensionally defined sets of SNOMED CT concepts (or expressions) that are required to test CDS rule criterion. Examples of CDS rules that use SNOMED CT expression constraints can be found in [3.1.3. Rule Examples](#).

For more information about expression constraints, please refer to [Expression Constraint Language - Specification and Guide](#).

Arden Syntax

The Arden Syntax is a widely-used and mature markup language for representing, sharing, and processing clinical knowledge, which makes it suitable in the application of expressing rules for use in decision support. The syntax has a long history, but is currently maintained by HL7 International. One of the advantages of the ARDEN syntax is improved human readability which is achieved by its resemblance to natural language. This in turn makes ARDEN code easier for non-technical audiences to interpret.

When used in CDS, Arden code can be embedded in independent files called medical logic modules (MLMs).^[1] The improved readability of Arden syntax makes it easier for a clinician to validate the clinical accuracy of any given MLM. MLMs have been widely used and libraries of these modules are available. It is also worth noting that the Arden syntax does not define how it should be integrated within an electronic health record or how an application should use it.

For more information on the Arden Syntax, please refer to the [HL7 Implementation Guide for Arden Syntax, Release 1](#).

FHIR CDS Resource

PlanDefinition is a general FHIR resource which can be used to represent a range of CDS artifacts such as rules, order sets, and protocols.^[2] According to the HL7 FHIR specification, a resource contains a set of structured data items that conform to the definition of the resource type and can be used to exchange and/or store data to satisfy a wide range of clinical and administrative healthcare information needs.^[3] PlanDefinition is currently defined as a draft resource within FHIR's [Clinical Reasoning](#) module. The Clinical Reasoning module is a draft of the Clinical Quality Framework Implementation Guide (or FHIR-Based Clinical Quality Framework). The guidance in this module is prepared as a Universal Realm Specification, which means it is designed to be used Internationally.

The PlanDefinition resource can be used to represent a rule using the Event-Condition-Action pattern. This pattern is defined within the actionDefinition element of the PlanDefinition resource. "A single, top-level actionDefinition represents the overall rule, with the triggerDefinition element used to specify the triggering event(s), the condition element used to specify the applicable condition for the rule, and the actionDefinition itself describing the action to be performed."^[4] The PlanDefinition resource is used to describe series, sequences, or groups of actions to be taken, while the ActivityDefinition resource is used to define each specific step or activity to be performed. An example of an XML instance of a PlanDefinition resource that encapsulates a Chlamydia Screening rule is shown below.

```

<PlanDefinition>
  <id value="chlamydia-screening-intervention"/>
  <identifier> ☒ </identifier>
  <version value="2.0.0"/>
  <title value="Chlamydia Screening CDS Example Using Common"/>
  <status value="draft"/>
  <description value="Chlamydia Screening CDS Example Using Common"/>
  <publicationDate value="2015-07-22"/>
  <topic> ☒ </topic>
  <library>
    <reference value="Library/ChlamydiaScreening_CDS_UsingCommon"/>
  </library>
  <actionDefinition>
    <title value="Patient has not had chlamydia screening within the recommended
timeframe..."/>
    <condition>
      <expression value="NoScreening"/>
    </condition>
    <dynamicValue>
      <path value="~/"/>
      <expression value="ChlamydiaScreeningRequest"/>
    </dynamicValue>
  </actionDefinition>
</PlanDefinition>

```

Figure 3.1.4-1: Example PlanDefinition XML instance

For more information on this FHIR resource, please refer to <http://build.fhir.org/plandefinition.html> or <http://build.fhir.org/clinicalreasoning-module.html>.

- 1 The Arden Syntax standard for clinical decision support: Experiences and directions, Samwald et. al
- 2 <http://build.fhir.org/plandefinition.html>
- 3 <https://www.hl7.org/fhir/resource.html>
- 4 <http://build.fhir.org/cqif/cqif-knowledge-artifact-representation.html#event-condition-action-rule>

3.2. Guidelines

The two main approaches to preparing clinical guidelines for use in CDS are to use simple markup (also referred to as "semantic tagging"), and to use a standard guideline representation language.

Simple guideline markup involves annotating free-text clinical guidelines using terminology codes that represent its meaning. This enables relevant guidelines to be retrieved based on specific codes recorded in a patient's health record. For more information on this approach, please refer to [3.2.1. Guidelines with SNOMED CT](#).

An alternative approach is to use a standard guideline representation languages to formally define each guideline. Some examples of standard guideline representation languages, such as the Guideline Definition Language (GDL) and the Guideline Interchange Format (GLIF) are presented in the section [3.2.2. Standards for CDS Guidelines](#).

3.2.1. Guidelines with SNOMED CT

This section examines how clinical guidelines can be linked to SNOMED CT to enable the automated display of contextually relevant knowledge resources. We begin by reviewing how a guideline can be linked to a SNOMED CT concept using semantic tags. Next we will examine how SNOMED CT concepts can be associated with guidelines using a reference set. Lastly, we will look at the automated display of a contextually relevant guideline, based on the selection of a SNOMED CT concept in a data entry protocol.

Linking Guidelines to SNOMED CT

One approach known as simple markup, involves the application of semantic tags using terminology codes (such as SNOMED CT concept identifiers) to free text clinical guidelines. When using SNOMED CT, concept identifiers (or expressions) are added as document metadata to the appropriate guideline or text within the guideline. For example, when applying semantic tags to asthma management guidelines, we might add the following concept identifiers, to enable the guideline to be linked to a patient's health record that includes a diagnosis of |Asthma|, a regime of |Asthma management|, or an assessment scale coded with |Asthma control questionnaire| (respectively):

Concept
195967001 Asthma (disorder)
406162001 Asthma management (regime/therapy)
445531003 Asthma control questionnaire (assessment scale)

The diagram below depicts the process of tagging the SNOMED CT concepts mentioned above to an Asthma Management Guideline.

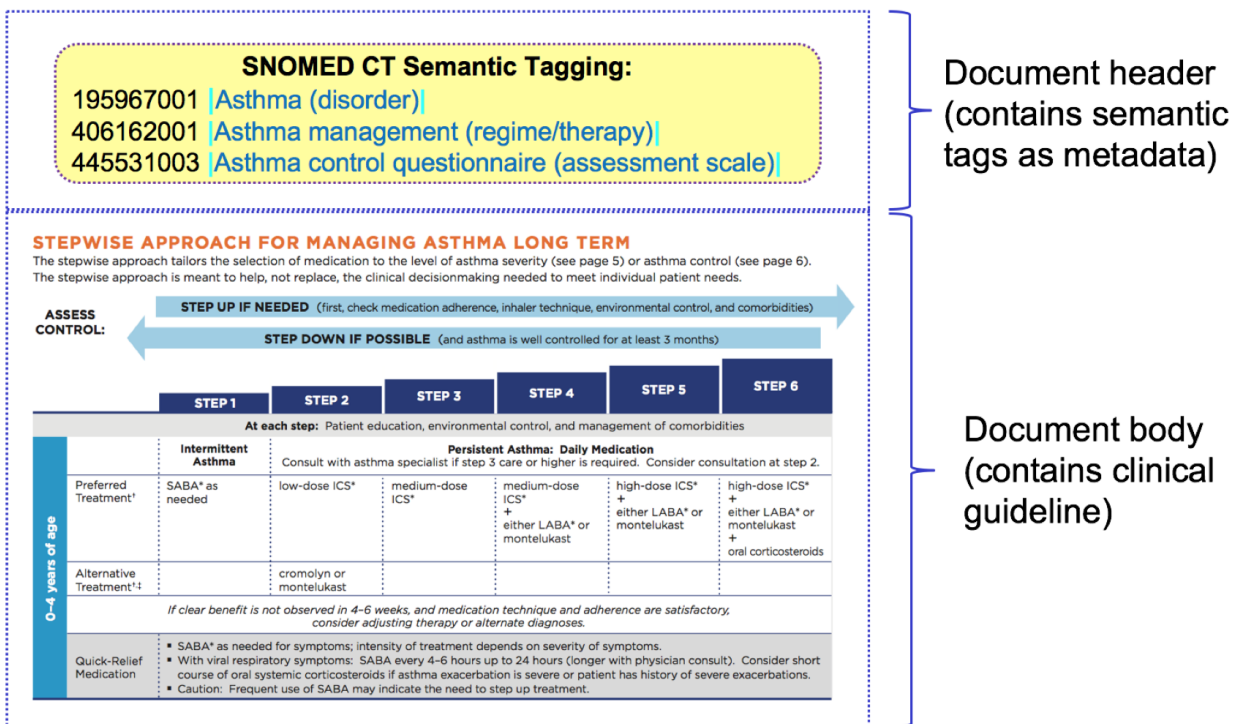


Figure 3.2.1-1: Asthma Management Guideline tagged with SNOMED CT concepts

This tagged resource may then be presented as reference material when a relevant clinical scenario arises. For example, the guidelines shown in Figure 3.2.1-1 could be presented upon the diagnosis of |Asthma|. Additional details on the mechanics of this process are provided below in the Selecting Relevant Guidelines section below.

Linking SNOMED CT to Guidelines

SNOMED CT Annotation Reference Sets can be used as a mechanism to define, share, and distribute links from SNOMED CT components to appropriate guidelines. This approach involves defining one or more links to relevant guidelines (using a URL) as a string based annotation for each relevant concept. An example of this is shown in Table 3.2.1-1. In this example, the same clinical guideline is relevant to more than one SNOMED CT concept.

Table 3.2.1-1: Concepts which refer to NIH: Asthma Care Quick Reference guideline

refsetId	referencedComponentId	Annotation
71999999107 Guideline annotation reference set	195967001 Asthma (disorder)	http://www.example.com/asthma_guideline
71999999107 Guideline annotation reference set	406162001 Asthma management (regime/therapy)	http://www.example.com/asthma_guideline
71999999107 Guideline annotation reference set	445531003 Asthma control questionnaire (assessment scale)	http://www.example.com/asthma_guideline

Another example use case is linking a specific clinical field, such as a diagnosis, to an appropriate clinical guideline. Given the more specific scope, it may be possible to avoid repeating the same guideline for multiple SNOMED CT concepts referenced by the reference set. Table 3.2.1-2 below, shows an example in which disorder concepts are linked to appropriate clinical guidelines.

Table 3.2.1-2: Respiratory diagnoses linked to relevant clinical guidelines

refsetId	referencedComponentId	Annotation
71999999107 Guideline annotation reference set	195967001 Asthma	http://www.example.com/asthma_guideline
71999999107 Guideline annotation reference set	32398004 Bronchitis	http://www.example.com/bronchitis_guideline
71999999107 Guideline annotation reference set	13645005 Chronic obstructive lung disease	http://www.example.com/COPD_guideline

Selecting Relevant Guidelines

EHR systems, designed to use clinical guidelines marked up with semantic tags, are able to display relevant guidelines to the user when a subtype (or self) of the semantic tag concept is recorded in the health record.

Below is a generic template that could be used in an EHR system to facilitate the display of an appropriate clinical guideline, when a relevant diagnosis is recorded.

```
IF diagnosis = << [| + $semanticTag|]
THEN display clinical guideline
```

For example, the above template could be used on a clinical guideline that has the semantic tag |Asthma|, generating the CDS rule:

```
IF diagnosis = << |Asthma|
THEN display NIH Asthma Care Quick Reference
```

Using the above CDS rule, if a clinician selects a diagnosis of |Chronic asthmatic bronchitis|, the NIH Asthma Care Quick Reference would be displayed because |Chronic asthmatic bronchitis| is a subtype of |Asthma|.

This enables contextually relevant clinical knowledge to be presented to the user, based on the specific codes recorded in the patient's health record. The diagram below illustrates this scenario:

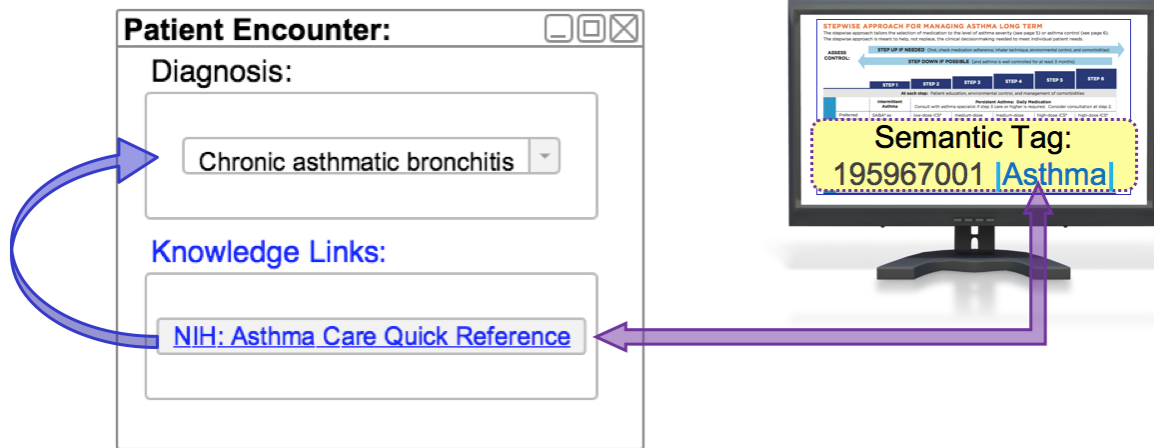



Figure 3.2.1-2: Upon entry of a specific diagnosis, a contextually relevant knowledge resource is presented on the user interface

 Asthma Care Quick Reference used for demonstration purposes.

3.2.2. Standards for CDS Guidelines

This section presents some examples of standards used to represent CDS guidelines. Please note that this list is not complete, and other standards and formalisms for representing clinical guidelines do exist.

Guideline Interchange Format

The Guideline Interchange Format (GLIF) is a language for modeling and executing clinical practice guidelines.¹ GLIF uses GELLO, and therefore can make use of SNOMED CT within its language. Users of GLIF have the option of viewing GLIF code in a interactive flowchart to represent guidelines and present pop-up information and instructions. The following screenshot presents the "UI view" of a query / guideline structured in GLIF. This example uses SNOMED CT to answer the questions in the decision shapes:²




Figure 3.2.2-1: UI view of a query / guideline structured in GLIF

For more information about GLIF, please refer to <https://kb.medical-objects.com.au/display/PUB/GLIF>.

Guideline Definition Language

The Guideline Definition Language (GDL)³ is a syntax designed to express clinical logic as rule inputs and outputs. Discrete rules can be combined together to support simple or complex decision making. The specification which accompanies the GDL describes its status as "trial." One of the goals of the GDL is to be able to share CDS artifacts across languages and technical platforms. GDL artifacts can be applied to point of care (POC) decision support and in population health analytics. The GDL uses reference and archetype models from openEHR and in doing so supports information model references that are language independent. For example, it is possible to add language translations without changing the logical definitions in the rule. It is also possible to bind locally defined terms in the guideline to a single concept, multiple concepts, or reference set in any reference terminology as the language considers a reference terminology to be an external resource.



For more information on the Guideline Definition Language, please refer to <http://www.openehr.org/releases/CDS/latest/docs/GDL/GDL.html>.

- 1 http://www.openclinical.org/gmm_glif.html
- 2 Content and diagram provided by Medical-Objects.
- 3 http://www.openehr.org/releases/CDS/latest/docs/GDL/GDL.html#_guideline_definition_language_gdl

3.3. Substrate

An important consideration in the development of a Clinical Decision Support System (CDSS) is the substrate over which the knowledge artifacts are authored and executed. When using SNOMED CT in a CDSS, the substrate is the SNOMED CT content over which the CDS rules are authored or executed. Because medical knowledge is constantly changing, it is important that the substrate over which CDS is applied is kept current. To support this requirement, SNOMED CT releases regular new versions of the terminology, and retains a history of changes using its strong versioning mechanism. With this in mind, both the [SNOMED CT edition](#) and the specific version of that edition (released on a given date) need to be considered in determining the SNOMED CT substrate. For more information on the topic of [versioning](#), please refer to section 11.4 in [Data Analytics with SNOMED CT](#).

Knowledge Artifact Substrate

When publishing CDS knowledge artifacts, such as rules or guidelines, it is important to clearly indicate the substrate over which the artifacts were authored. The substrate used to author a CDS artifact needs to be considered when determining the appropriate substrate to use in the execution of that artifact.

For example a CDS rule, using the SNOMED CT US edition, dated 20160901 (September 1st 2016), may refer to the concept [|Persistent asthma|](#). If this rule was executed against the SNOMED CT International Edition (20170131), then this extension concept would not be found, and the rule could not be executed.

Similarly, a CDS rule using the International edition (20170131) may refer to the concept [|Dual energy computed tomography|](#). If the rule was executed against the 20160731 International edition (or an older version), then this concept (created in the 20170131 version) would not be found, and the rule could not be executed. This would also be true if the same CDS rule was executed against the US edition (20160901), because this edition is dependent on the 20160731 International edition.

Therefore, CDS knowledge artifacts referencing SNOMED CT concepts must be executed (by the inference engine) using a SNOMED CT substrate that includes the same [modules](#), or a superset of the modules included in the SNOMED CT substrate used to author the artifact. In addition, the SNOMED CT substrate used to execute the CDS artifacts must use the same version, or a more recent version of these SNOMED CT modules, to ensure that all the referenced concepts are present. Please note that if a newer version of the substrate is used to execute the rules, then it is possible that a concept or relationship used at the time of artifact authoring may have become inactive. Edition and version dependencies such as these should be checked when adopting new CDS artifacts or updating a CDS system to use a newer version of SNOMED CT.

Electronic Health Record Substrate

Another consideration when selecting the SNOMED CT substrate on which to execute CDS artifacts, is the substrate used to record data in the Electronic Health Record (EHR).

For example, if a CDS rule is triggered when the diagnosis recorded in the EHR is a descendant of [|Asthma|](#), that is:

```
IF < |Asthma| THEN ...
```

and this rule is executed over the SNOMED CT International edition, dated 20170131 (January 31st 2017), then EHR records which capture a diagnosis of [|Persistent asthma|](#) from the SNOMED CT US edition (20160901) will be unsuccessful in triggering the CDS rule.

Similarly, if a CDS rule is triggered when the diagnosis recorded in the EHR is a descendant of [|Pulmonary disease|](#), that is:

```
IF < |Pulmonary disease| THEN ...
```

and this rule is executed over the US edition (20160901), then EHR records which capture a diagnosis of [|Squamous cell carcinoma of left lung|](#) will be unsuccessful in triggering the CDS rule. This is because the US edition (20160901) is dependent on the International edition (20160731), and the concept referenced above was added to the International edition (20170131).

Therefore, CDS knowledge artifacts must be executed (by the inference engine) using a SNOMED CT substrate that includes the same modules, or a superset of the modules used by the EHR system to record patient data. In addition, the SNOMED CT substrate used to execute the CDS knowledge artifacts must use the same or a more recent version of these SNOMED CT modules, to ensure that all the referenced concepts are present. Please note that if a newer version of the substrate is used to execute the rules, then it is possible that a concept recorded in the EHR may have become inactive. Edition and version dependencies such as these should be checked when implementing new CDS rules over existing EHR data or updating a CDS system to use a newer version of SNOMED CT.

4. Inference Engine

At the heart of a clinical decision support system, the inference engine uses inputs from the user, the record services, and the terminology services to process the machine readable rules, guidelines, or CDS artifacts. It is the job of the inference engine to establish if the CDS conditions have been met and determine the appropriate outcome. It does this by executing queries over the health records and terminology, to test the CDS conditions defined in the CDS rules. Note that it is the **communications mechanism** which handles the action defined in the CDS rules, but the inference engine determines whether or not the action should be carried out.

The diagram below illustrates the key inference engine interactions described above:

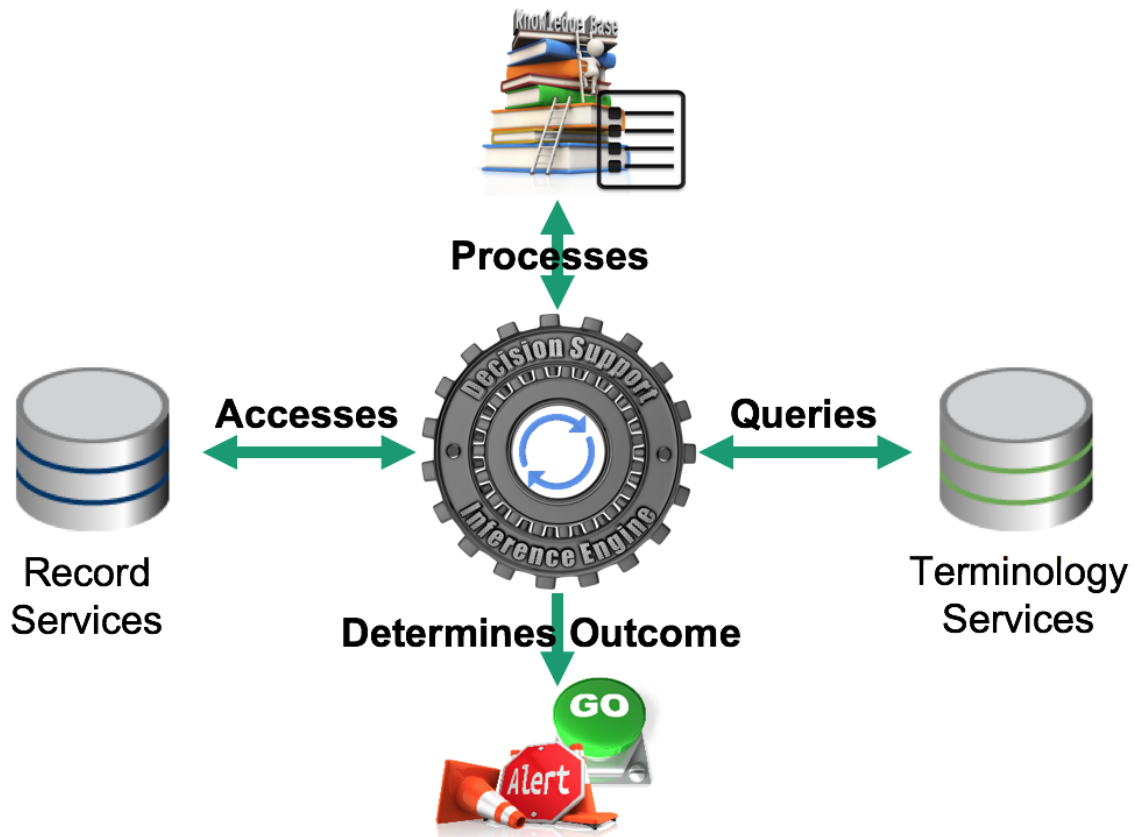


Figure 4-1: Inference engine key interactions

The following topics, which relate to the inference engine, are explored in the following sections:

- 4.1. Reasoning with SNOMED CT
- 4.2. Accessing Clinical Records
- 4.3. Accessing Terminology



4.1. Reasoning with SNOMED CT

Features of SNOMED CT can be used in a range of techniques which may then be applied to clinical decision support. For example, these techniques can help to execute decision support logic by assisting the inference engine in evaluating the trigger conditions defined in CDS rules.

This section describes these SNOMED CT techniques with respect to CDS by first providing an overview of the technique, and then presenting an example of how the inference engine can apply the technique to execute a specific CDS rule.

The following SNOMED CT techniques can be used by the CDS inference engine:

- [4.1.1. Reasoning with Subsets](#)
 - [4.1.2. Reasoning using Subsumption](#)
 - [4.1.3. Reasoning using Defining Relationships](#)
 - [4.1.4. Reasoning with Description Logic](#)
-

4.1.1. Reasoning with Subsets

Overview

A subset is defined in mathematics as a set whose members are all contained in another set. A SNOMED CT Subset typically refers to a collection of components that all come from the same edition of SNOMED CT. This is depicted in the diagram below.

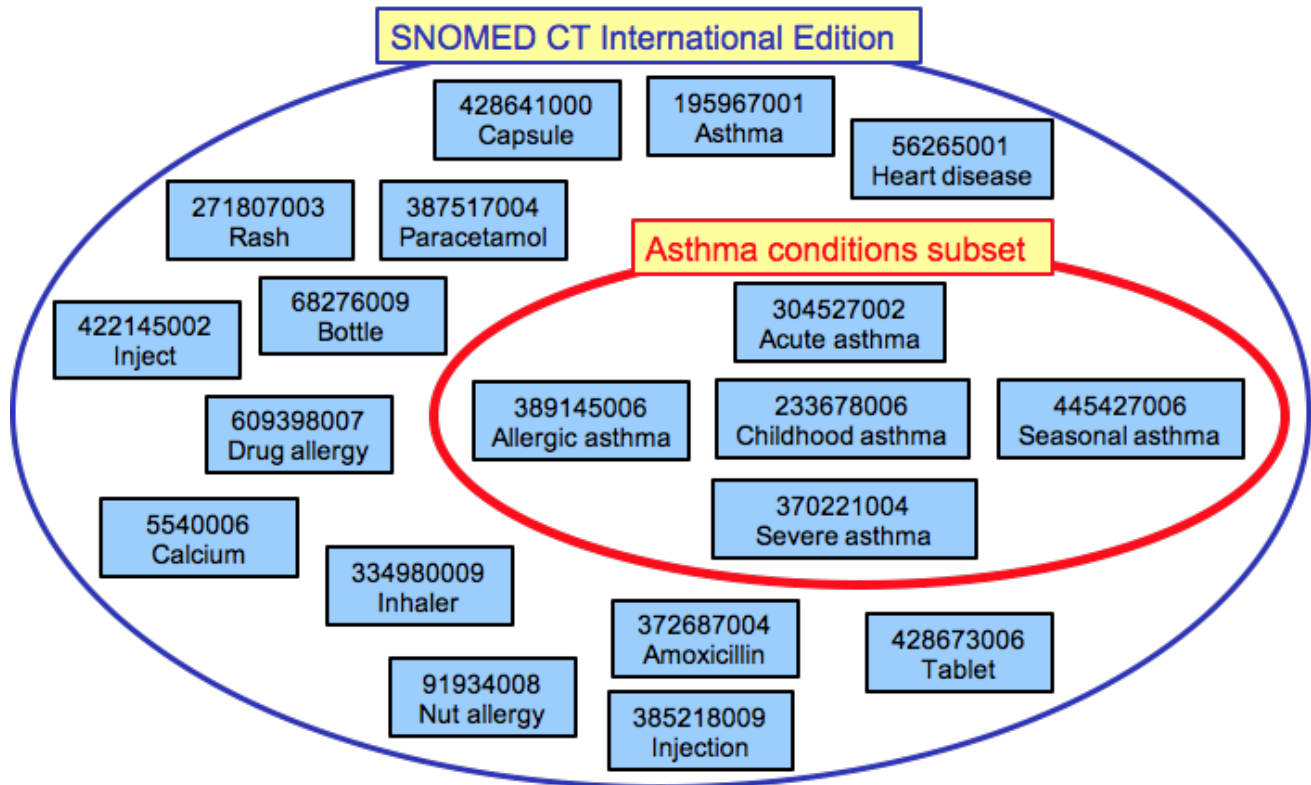


Figure 4.1.1-1: A subset of concepts related to the diagnosis of asthma is selected from the International Edition of SNOMED CT

A SNOMED CT subset may be defined extensionally, by enumerating all of the components in the set or intensionally, by defining a query written using the Expression Constraint Language - Specification and Guide.

Extensionally and intensionally defined subsets can both be represented as SNOMED CT reference sets, which support versioning and traceability. For more information about reference sets, please refer to the [Practical Guide to Reference Sets](#). For additional information on using subsets in queries, please refer to [6.1 Subsets in Data Analytics with SNOMED CT](#).

Example

This section presents a simple example of a CDS rule defined using a SNOMED CT subset, and explains how this rule could be executed by the CDS inference engine.

CDS Rule

The diagram below shows a simple CDS rule based on the IF-condition-THEN-action pattern. This rule uses a SNOMED CT subset to define the set of diagnoses that should trigger the display of the asthma management guidelines. It can be read as follows - "IF the diagnosis is a member of the Asthma conditions reference set THEN display the asthma management guidelines".



Figure 4.1.1-2: CDS rule which uses fictitious "Asthma conditions ref subset" in its definition

Execution of Rule

When executing this rule, the inference engine checks the given diagnosis for membership in the [239999999106 |Asthma conditions reference set|](#). The associated SNOMED CT subset is defined extensionally using a simple type reference set, and its members can be queried using a standard SNOMED CT terminology service.

The diagram below illustrates the process followed by the inference engine in executing the CDS condition in the above rule, when the clinician selects a diagnosis of [|Occasional asthma|](#). The inference engine checks if this concept is a member of the [239999999106 |Asthma conditions reference set|](#), and determines that it is not a member. As a result, the condition evaluates to false, and the action is not triggered.

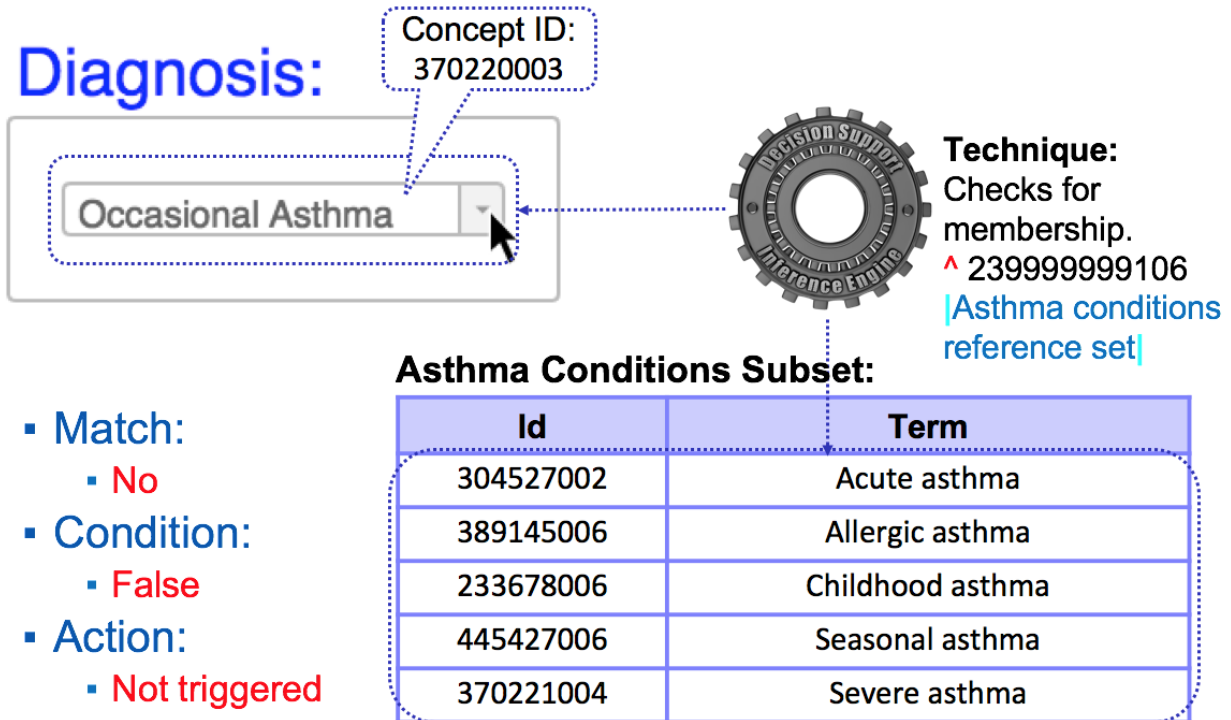


Figure 4.1.1-3: The inference engine compares the diagnosis entered against a predefined Asthma Conditions Subset

4.1.2. Reasoning using Subsumption

Overview

One of the fundamental benefits of SNOMED CT is its built-in polyhierarchy that specifies which concepts are subtypes of others. This hierarchy facilitates the automated grouping of health records which have been encoded using SNOMED CT. The **|is a|** relationships in SNOMED CT form the basis of its subtype hierarchy.

For example, **|Fracture of shaft of femur|** has an **|is a|** relationship to **|Fracture of femur|**, and therefore (as the diagram below illustrates), the concept **|Fracture of shaft of femur|** is subsumed by **|Fracture of femur|**. ¹

Figure 4.1.2-1: Example of subsumption

This also means that if a patient has a **|Fracture of shaft of femur|**, then it is implied (i.e. it is also true) that they have a **|Fracture of femur|**. We can use this principal to aggregate health records that have been encoded with SNOMED CT. By selecting any code that is a subtype of **|Fracture of femur|**, we are selecting all the codes that imply that **|Fracture of femur|** is true (given the appropriate context).

When testing for subsumption, we must also consider the transitivity of the **|is a|** relationship. For example, the diagram below indicates that **|Severe persistent asthma|** is a subtype of **|Severe asthma|** which is a subtype of **|Asthma|**. Therefore **|Severe persistent asthma|** is also a subtype of **|Asthma|**.




Figure 4.1.2-2: Example of subsumption and transitivity

As previously suggested in the section [1.4. SNOMED CT Features](#), the hierarchical relationships of SNOMED CT can be leveraged to enable clinical decision support. More specifically, we can apply subsumption testing to make additional determinations. For additional information on subsumption, please refer to [6.2 Subsumption](#) in [Data Analytics with SNOMED CT](#).

Example

CDS Rule

The diagram below shows a simple CDS rule based on the IF-condition-THEN-action pattern. This rule uses the descendant or self operator (**<<**) from the [Expression Constraint Language - Specification and Guide](#) to check if the diagnosis is in the set of concepts that includes **|Asthma|** and all of its subtypes.



Figure 4.1.2-3: CDS rule defined using subsumption

Execution of Rule

When executing this rule, the inference engine tests if the given diagnosis is subsumed by the concept **|Asthma|**. This subsumption testing can be performed using a range of approaches, including using a [7.5.2 Transitive closure implementation](#). A transitive closure table facilitates rapid testing of all possible **|is a|** relationships, and provides a very effective way of testing concept subsumption in relational databases.

The diagram below illustrates the process followed by the inference engine in executing the CDS condition in the above rule, when the clinician selects a diagnosis of **|Mild persistent asthma|**. The inference engine checks if this concept is a subtype of **|Asthma|**, and determines that it is. As a result, the condition evaluates to true, and the action is triggered.

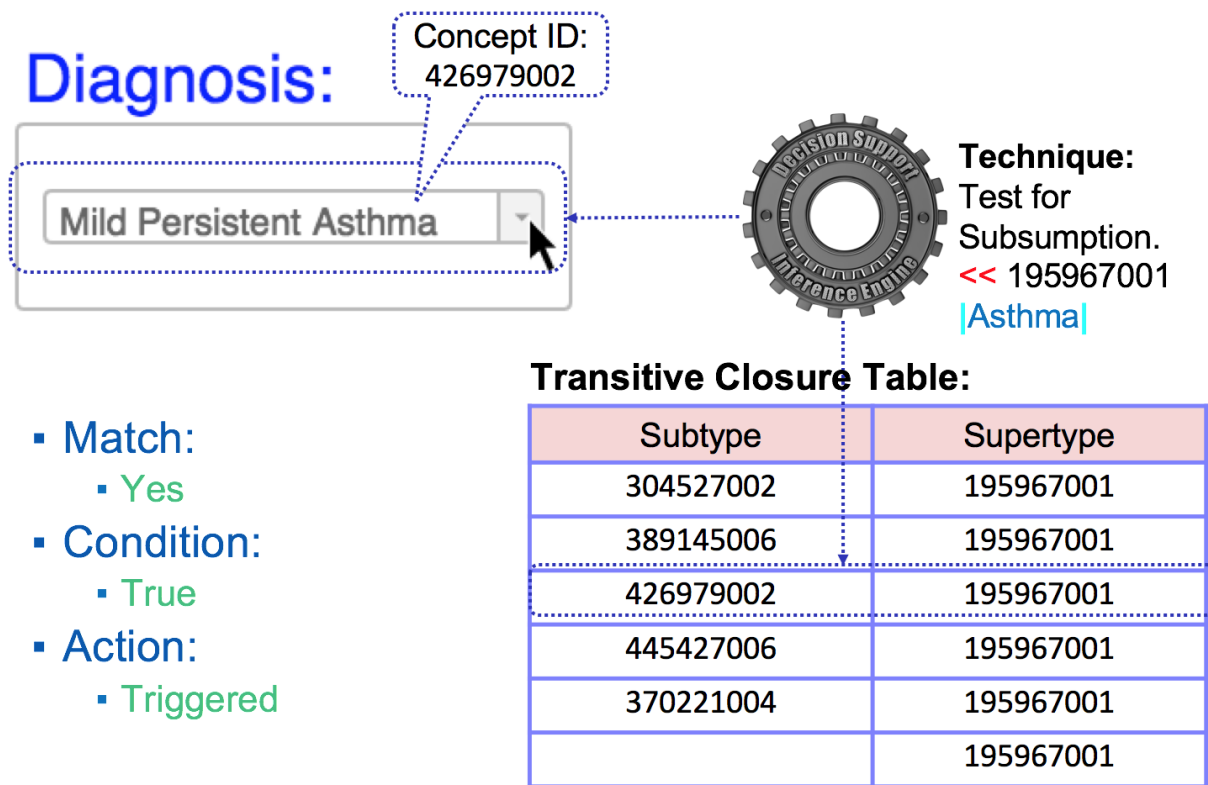


Figure 4.1.2-4: The inference engine checks if the diagnosis entered is a subtype of |Asthma|

Diagram uses notation from Diagramming Guideline.

4.1.3. Reasoning using Defining Relationships

Overview

In addition to the subtype relationships in SNOMED CT, attribute relationships may be used to support the definition of concepts. Only the relationships that are necessary (i.e. always true) are used as defining relationships in SNOMED CT. This is because these are the ones that produce reliable and consistent inferences. For example:




Figure 4.1.3-1: Concept definition consisting of subtype (blue arrows) and defining (green arrows) relationships

The **green** arrows in the diagram above show that the concept |myocardial infarction| has two necessary attribute relationships that represent a characteristic of the meaning of the concept. It always has an |associated morphology| of |infarct|, and it always has a |finding site| of |myocardium structure|. The **blue** arrows in the diagram above are used to indicate the subtype relationships.

The full definition of a concept consists of both the defining subtype relationships and the defining attribute relationships. There are over 50 attributes in SNOMED CT which can each be used as the "type" of a defining relationship, including |causative agent|, |method|, and |laterality|.

The **SNOMED CT Concept Model** provides rules about how these attributes can be used to define concepts from different hierarchies. The **SNOMED CT Machine Readable Concept Model (MRCM)** represents these rules in a form that can be read by a computer and applied to test that CDS criterion comply with these rules.

As previously suggested in section 1.4. **SNOMED CT Features**, the defining relationships of SNOMED CT can be leveraged to support CDS. For additional information on using SNOMED CT defining relationships in queries, please refer to 6.3 **Using Defining Relationships of Data Analytics with SNOMED CT**.

Example

CDS Rule

The diagram below shows a simple CDS rule based on the IF-condition-THEN-action pattern. This rule uses attribute refinements in the **SNOMED CT Expression Constraint Language** to define the set of procedures with a |Procedure site| that is a type of |Structure of the respiratory system|.




Figure 4.1.3-2: CDS rule which uses a defining relationship in its definition

Using attribute refinements in the CDS rule criteria facilitates richer expressivity and specificity in the rules. For example, we can restrict pharmaceutical/biological products based on their active ingredients, procedures based on their methods, and disorders based on their finding sites.

Execution of Rule

When executing this rule, the inference engine must process the defining relationships of each |Procedure| concept to determine which ones have a |Procedure site| that is a type of |Structure of respiratory system|. These relationships are distributed as part of SNOMED CT's Release Format 2 relationship file, which can be searched by a terminology service to discover relationships that match the given attribute refinement.

The diagram below illustrates the process followed by the inference engine in executing the CDS condition in the above rule, when the clinician selects the therapy |Intermittent continuous positive airway pressure|. Once the inference engine has found the defining relationships whose source is |Intermittent continuous positive airway pressure| and whose type is |Procedure site|, it determines whether the destination of these relationships is either |Structure of respiratory system| or a subtype of |Structure of respiratory system| (e.g. using a transitive closure table). Since the given procedure has a |Procedure site| equal to |Structure of respiratory system|, the condition in the rule evaluates to true, and the action is triggered.

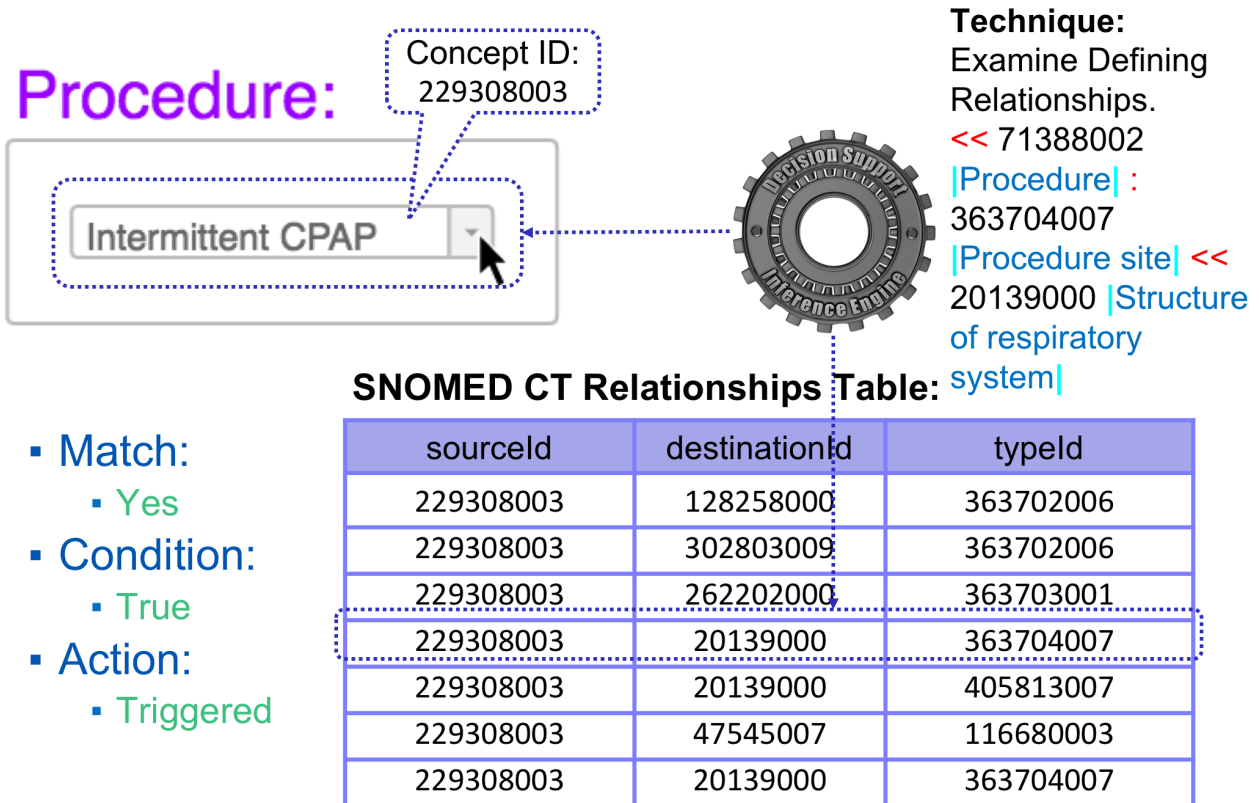


Figure 4.1.3-3: The inference engine checks if the diagnosis entered has a defining relationship which states that the |Procedure site| is |Structure of respiratory system| or a subtype.

4.1.4. Reasoning with Description Logic

Overview

Description logic (DL) reasoners can apply additional logic-based techniques to assist with clinical decision support reasoning. Two of the DL techniques supported by SNOMED CT have been briefly described below.

It is also worth pointing out that DL can be applied over the terminology or to the terminology in combination with the record structure. For more information on this topic, please refer to sections [6.4 Description Logic Over Terminology](#) and [6.5 Description Logic Over Terminology and Structure](#) in the guide [Data Analytics with SNOMED CT](#).


Expression Subsumption

SNOMED CT supports the use of postcoordinated expressions to define additional clinical meanings beyond the standard precoordinated concepts. Postcoordinated expressions are comprised of two or more concepts, and are structured in accordance with the [Compositiona l grammar](#). A simple example of a postcoordinated expression is:

```
399963005 |Abrasion| :
  363698007 |Finding site| = 67269001 |Skin structure of ankle|
```

This can be read as "an abrasion with a finding site of skin of ankle".

When postcoordinated expressions are used to capture and record clinical meaning in a health record, the CDS inference engine may need to be able to test if one expression subsumes another to execute the CDS rules. For example, it would be reasonable to conclude that the first expression listed below, subsumes the second expression if you were aware that [|Mild pain|](#) is subsumed by [|Pain|](#).



However, many expressions are much more complex than this and may involve multiple focus concepts, attribute groups, and nesting as examples. There are a couple of methods which can be utilized to determine if one expression subsumes another. The first process, which is a manual process, involves [12.4.2 Normalize Expression](#) the expressions, comparing the primitive focus concepts, and then comparing the defining attributes. The other is an automated process which can be used to compare expressions for subsumption. Expressions can be imported into a description logic classifier and classified in the same way as SNOMED CT concept definitions, which tests for subsumption in the process.

Property Chaining

The transitive nature of the [|is a|](#) attribute allows us to make subtype inferences. This [topic](#) was explored in the section on [Reasoning using Subsumption](#). In some cases, different types of attribute relationships may be related to one another in such a way that additional inferences are possible. For example (using fictitious concepts) if Lucy [|Has sister|](#) Beth, and Beth [|Has daughter|](#) Jane, then Lucy [|Has niece|](#) Jane. The chain from [|Has sister|](#) to [|Has daughter|](#) implies [|Has niece|](#). This rule can be expressed as:

```
|Has sister| o |Has daughter| = |Has niece|
```

At present, the only property chain recognized in the International Edition of SNOMED CT is from [|Direct substance|](#) to [|Active ingredient|](#) and can be expressed as such:

```
|Direct substance| o |Has active ingredient| = |Direct substance|
```

This can be used to provide a link from product administration (as part of a procedure) to substance administration. Additional property chains can be added at the local implementation level, if required.

4.2. Accessing Clinical Records

This section describes the approaches that inference engines use to access EHR records for CDS. The CDS rules that an inference engine executes typically include references to EHR records and terminology. There are two general approaches to accessing health records from these CDS artifacts. The first is the direct access approach, in which a reference to the physical store is used. A simple example of this would be a reference to a patient diagnosis in a database table. The other approach is to base the CDS references on a common logical information model, and then map this to one or more physical datastores, as required. This approach enables a more standardized approach to the development of CDS rules and other artifacts. These two approaches are described in more detail below, along with some of the advantages, challenges, and examples. Please note that standards for accessing clinical records will be discussed in section 4.2.1. [Standards for Accessing Clinical Records](#). Approaches to access the terminology which will be discussed in section 4.3. [Accessing Terminology](#).

Direct Access

Perhaps the most obvious approach to accessing health records is to use direct references to the locations in the clinical record store that will be used in a CDS rule. Many consider this the more common approach for point of care (POC) decision support today. For example, the pointers within a CDS rule could be expressed to reference a schema, table, and column in an SQL database. This approach requires a detailed knowledge of how the EHR data is stored to achieve an appropriate outcome from CDS tools. The challenge with this approach is that it can be more difficult to share CDS artifacts across institutions that may use different physical data stores.

Logical Model

The logical model approach aims to standardize all references to EHR data, by specifying a common logical information model upon which all CDS rules are defined. This enables CDS rules to be shared and reused in different physical implementations. However, an additional transformation is usually required for each physical EHR store, to convert the logical references into physical ones. Examples of the logical model approach are presented in section 4.2.1. [Standards for Accessing Clinical Records](#).

4.2.1. Standards for Accessing Clinical Records

This section presents some examples of standards for accessing clinical records. Please note that this list is not complete, and other standards and formalisms for accessing clinical records do exist.

Clinical Information Modeling Initiative

The HL7 Clinical Information Modeling Initiative (CIMI) is an HL7 International working group, which aims to improve the interoperability of healthcare systems by providing a shared open library of implementable clinical information models^[1]. CIMI clinical models, which are defined using computable formalisms such as the Archetype Definition Language (ADL) and Archetype Modeling Language (AML), are based on a common reference model using a common set of data types. CIMI models also have formal bindings to standard terminologies, including SNOMED CT and LOINC. SNOMED CT has been selected as the primary reference terminology for CIMI's clinical models. A number of CDS efforts within HL7 International are expected to use CIMI clinical models as the basis for referencing clinical data within CDS artifacts. For more information on CIMI please refer to [Clinical Information Modeling Initiative \(opencimi.org\)](http://opencimi.org) and [Clinical Information Modeling Initiative \(HL7 work group\)](#).

Quality Information and Clinical Knowledge model

The Quality Improvement Core (QICore) Implementation Guide is a U.S. realm-specific CDS initiative that references a logical model called the Quality Information and Clinical Knowledge (QUICK) model. The QUICK model (which is expected to be aligned with CIMI formalisms) will provide a uniform way for clinical decision support and quality measures in the U.S. to refer to clinical data^[2]. The QUICK logical model is defined as a series of QICore specific FHIR profiles. It provides a way for applications to access data using FHIR interfaces. Several of these QICore profiles have bindings to SNOMED CT value sets in their definition. For example, the Condition model (shown below) binds a 'SNOMED CT Body Structure' value set to the data element 'bodySite'.

Condition			
Key: ✓ = Must support, ★ = QICore-defined extension, ! = Is-Modifier			
Field	Card.	Type and Description	doco
✓ abatement[x]	0..1	DateTime Age Boolean interval<DateTime> interval<Quantity> String The date or estimated date that the condition resolved or went into remission. This is called "abatement" because of the many overloaded connotations associated with "remission" or "resolution" - Conditions are never really resolved, but they can abate.	
asserter	0..1	Patient Practitioner Individual who is making the condition statement.	
✓ bodySite	0..*	List<Concept> The anatomical location where this condition manifests itself. Binding: SNOMED CT Body Structures (Example)	
✓★ bodySiteInstance	0..*	List<BodySite> An Extension	

Figure 4.2.1-1: QUICK Conditions model

The QUICK logical model provides the basis upon which the FHIR RESTful interfaces refer to clinical data in a CDS service. For more information on QICore please refer to [Quality Improvement Core \(QI-Core\) Implementation Guide](#).

Virtual Medical Record

The Virtual Medical Record (vMR) for CDS is an HL7 standard, which describes a standardized "virtual interface" for CDSSs to refer to the data in clinical records. The vMR logical data model is based on the HL7 V3 RIM (Reference Information Model). Developing CDS rules can be time-consuming and costly. Hence, the key goal of the vMR is to provide a simplified common information model upon which sharable CDS resources can be developed^[3]. To implement the vMR, each EHR system must create a virtual interface that exposes its clinical data in the standardized vMR format, to facilitate shared CDS logic working across multiple EHR systems.^[4] An example of an expression, written in terms of the HL7 vMR, which could be used in a CDS rule, is shown below. This expression asserts that a patient has a condition of |Asthma| that has a status of |Active|.

Example from HL7 Version 3 Standard: Clinical Decision Support; Virtual Medical Record (vMR) Logical Model, Release 2

```

clinicalstatement[xsi:type="vmr:Problem" and
/templateId[root="2.16.840.1.113883.3.1829.11.7.2.5"] and
/conditionCode[codeSystem="2.16.840.1.113883.6.96" and code="195967001"] and
/conditionEffectiveTime[/low[value<="20130814"]]] and
/conditionStatus[codeSystem="2.16.840.1.113883.6.96" and code="55561003"]
]

```

Figure 4.2.1-2: vMR expression example

The vMR's data model includes clinical findings, problems, allergies, adverse events and patient history. The vMR was also optimized to permit CDS languages such as GELLO to reference a standard model of clinical data. For more information about the HL7 vMR please refer to the HL7 Version 3 Standard: Clinical Decision Support; Virtual Medical Record (vMR) Logical Model, Release 2.

GELLO

GELLO⁵ is an object-oriented programming language that can be used to support access to health record data in CDS. It has been adopted by ANSI and HL7 as a language used in CDS and GELLO Release 2 now part of the HL7 v3 product suite.

GELLO provides a standardized interface and query language for accessing data in health information systems. Expressions can also be defined to compare data values and attributes. These values and attributes can then be used in decision support knowledge resources such as rules and guidelines. GELLO works hand in hand with the *Virtual Medical Record (vMR)*. A major advantage of this approach is that GELLO code can be used in different environments where health data is stored using a variety of formats and technologies.

Using GELLO with the vMR ensures that the code does not alter the physical medical record. It can also be used to answer complex queries and to query a reference terminology such as SNOMED CT. The example screen shot below illustrates how SNOMED CT refinements can be used in a GELLO expression:

```

1 imports SnomedCT
2
3 • Let utterances:Sequence(SnomedUtterance) =
4   Sequence(
5     SnomedUtterance{code='31978002'},
6     SnomedUtterance{code='65363002'},
7     SnomedUtterance{code='46866001 |Fracture of lower limb (disorder)| :'+
8       ' { 363698007 |Finding site (attribute)| = 12611008 |Bone structure of tibia (body structure)|, '+
9         ' 116676008 |Associated morphology (attribute)| = 72704001 |Fracture (morphologic abnormality)| }
10    }
11
12 • utterances->collect(u)
13 Tuple(
14   utterance=u,
15   validate=u.validate(),
16   canonical_form_short = u.canonical_form(false),
17   preferred_name=u.preferred_name(),
18   top_level_concept_id=u.top_level_concept_id(),
19   verbose_form=u.verbose_form(),
20   synonyms=u.synonyms()
21 )

```

Figure 4.2.1-3: GELLO expression using SNOMED CT refinement

For more information about GELLO, please refer to HL7 Version 3 Standard: GELLO, A Common Expression Language, Release 2 or <http://gello.org>.

¹ <http://www.hl7.org/Special/Committees/cimi/overview.cfm>

² <http://hl7.org/fhir/us/qicore/2016Sep/index.html>

³ HL7 Decision Support Service (DSS) and Virtual Medical Record (vMR) Standards, and OpenCDS Open-Source Implementation presentation, Kawamoto

⁴ http://www.hl7.org/implement/standards/product_brief.cfm?product_id=338

⁵ <https://kb.medical-objects.com.au/display/PUB/GELLO>

4.3. Accessing Terminology

As references to terminology codes may exist in both health records and CDS artifacts, the inference engine needs to be able to access terminology services to execute CDS logic. Furthermore, to maximize the benefits of using SNOMED CT, additional terminology operations, such as finding descendants of a concept and finding the value of a defining relationship, can be used. For example, a CDS rule may refer to all the descendants of |Heart disease| to ensure that a specific cardiology CDS rule is applied to all applicable diagnoses. Similarly, a CDS rule may need to find all the active ingredients in a medication, to ensure that a contra-indication does not occur.

This section discusses some of the options that a CDS inference engine can use to access terminology content.

Terminology Services

As discussed in section 2.1. *EHR System Architecture*, terminology services are those services required to load, update, access and make effective use of terminology content. Terminology services provide important functions to CDSSs, such as term searching (i.e. synonyms), definitional and reference set querying, and retrieval of map data. It is also useful if Terminology Services support the execution of the *Expression Constraint Language*, as this is a standardized way of representing terminology queries in CDS rules. For more information on Terminology services, please refer to the *SNOMED CT Terminology Services Guide*.

SNOMED CT APIs

An application programming interface (API) for a SNOMED CT enabled terminology server can be used to execute SNOMED CT searches and queries. The principle benefit of using a terminology server API is the reusability. Other systems are able to access terminology services without having to re-implement their functionality. Another key benefit is that the internal workings of the solution can be modified, improved, upgraded without impacting the external interfaces. For example, SNOMED CT can be updated, without necessitating any changes to the external systems which use terminology services. A number of commercial terminology servers offer proprietary APIs that enable SNOMED CT search and query. These include Dataline's SnAPI solution and B2i's Snow Owl Terminology Server. ¹ The following diagram depicts how the individual terminology services interact with the terminology store:

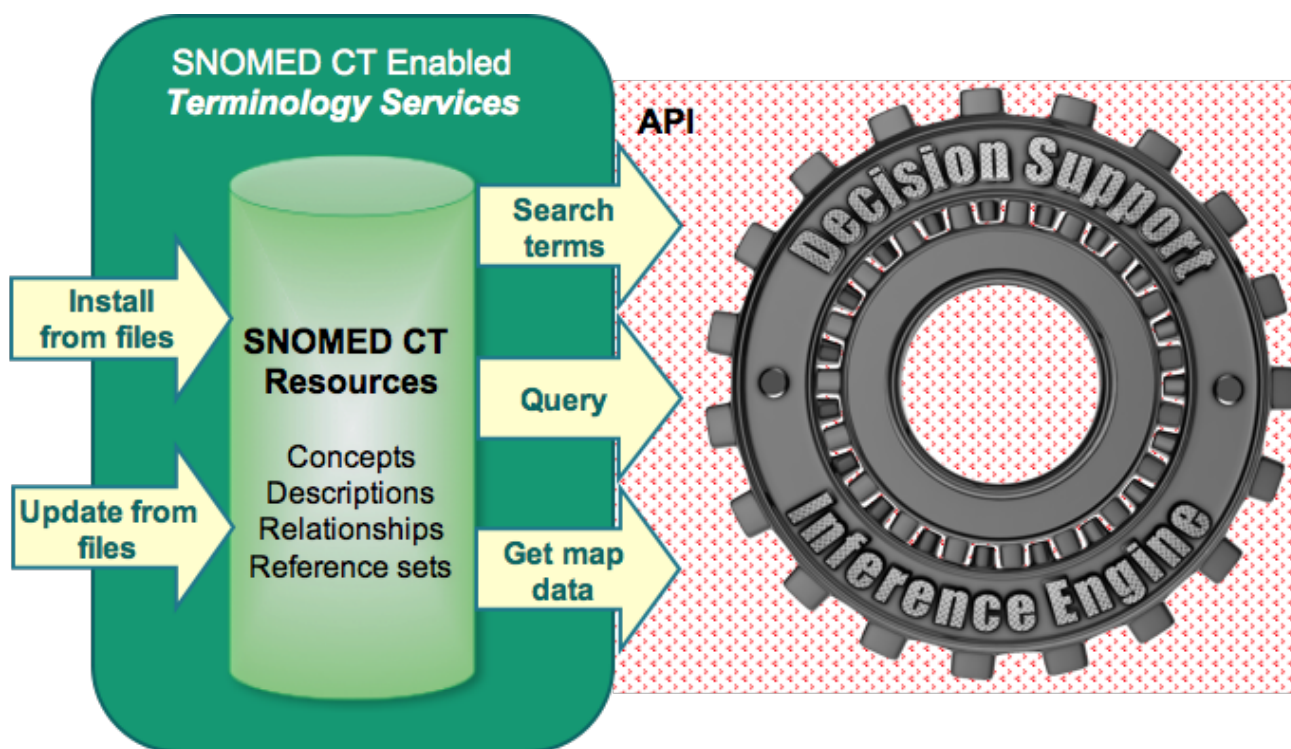


Figure 4.3-1: Terminology services and terminology store interactions

Services that load the terminology data into the server, either for installation or updating are illustrated on the left, while services which search and query over the installed terminology content are depicted on the right. The diagram also shows how the services depicted on the right could be made available to other services and components such as a CDS inference engine through the use of an API.

Standardized Terminology APIs

Standardized APIs for terminology services are also available. For example, HL7's specification for a FHIR Terminology Service, which is described as a service that lets healthcare applications make use of codes, code systems, and value sets without having to become experts in the fine details of terminology. ² The services provided include code lookup and validation, value set expansion, subsumption



testing, and maintaining a transitive closure table. HL7 has also published Common Terminology Services 2 (CTS2) which provides a standardized API that supports access to terminology servers which may contain a variety of code systems, including SNOMED CT.

[1](http://snomed.org/analytics) <http://snomed.org/analytics>

[2](http://hl7.org/fhir/2016Sep/terminology-service.html) <http://hl7.org/fhir/2016Sep/terminology-service.html>

5. Communications

Overview

The purpose of the communications mechanism is to handle CDS communications into and out of the system. Examples of user inputs include entry of clinical data, and the selection of a proposed drug, order set, or treatment regime. Examples of outputs include CDS interventions such as alerts, guidelines, diagnostic refinements, and smart forms. These outputs are typically delivered to the user interface. SNOMED CT has limited involvement in the communications mechanism of CDS as most of the codes and features will be used by the knowledge base and inference engine. That being said, it is possible that SNOMED CT terms are used at the user interface level as part of the data entry process. For more information on using SNOMED CT to support data entry, please refer to the [Search and Data Entry Guide](#). SNOMED CT can also be used in the CDSS outputs. For example, using the relevant terms in the alert messages, populating smart forms with SNOMED CT codes, or linking terms in CDS guidelines to other appropriate clinical knowledge sources.

The figure below depicts the key interactions of the CDS communications mechanism.




Figure 5-1: Communications key interactions

Once the inference engine has determined that an intervention is appropriate, the communications mechanism takes over and handles its delivery. Conversely, user inputs are also delivered into the CDSS by the communications mechanism. Note that guidelines or knowledge resources may reference externally hosted content, which may be accessed by the user via a link. An example of this would be a [PubMed](#) citation for biomedical literature.

Note that the diagram also shows how the internal CDSS communications (associated with the external inputs and outputs) are related to the components of the CDS rule. The communication 'inputs' feed into the event (from "ON event") and the condition (from "IF

condition") components of the rules, while the 'outputs' are the result of the action (from "THEN action") that is performed if the event occurs and condition is true.

Example

The following screenshot^[2] was generated from an EHR with CDS capabilities. This illustrates what a typical CDS intervention may look like.

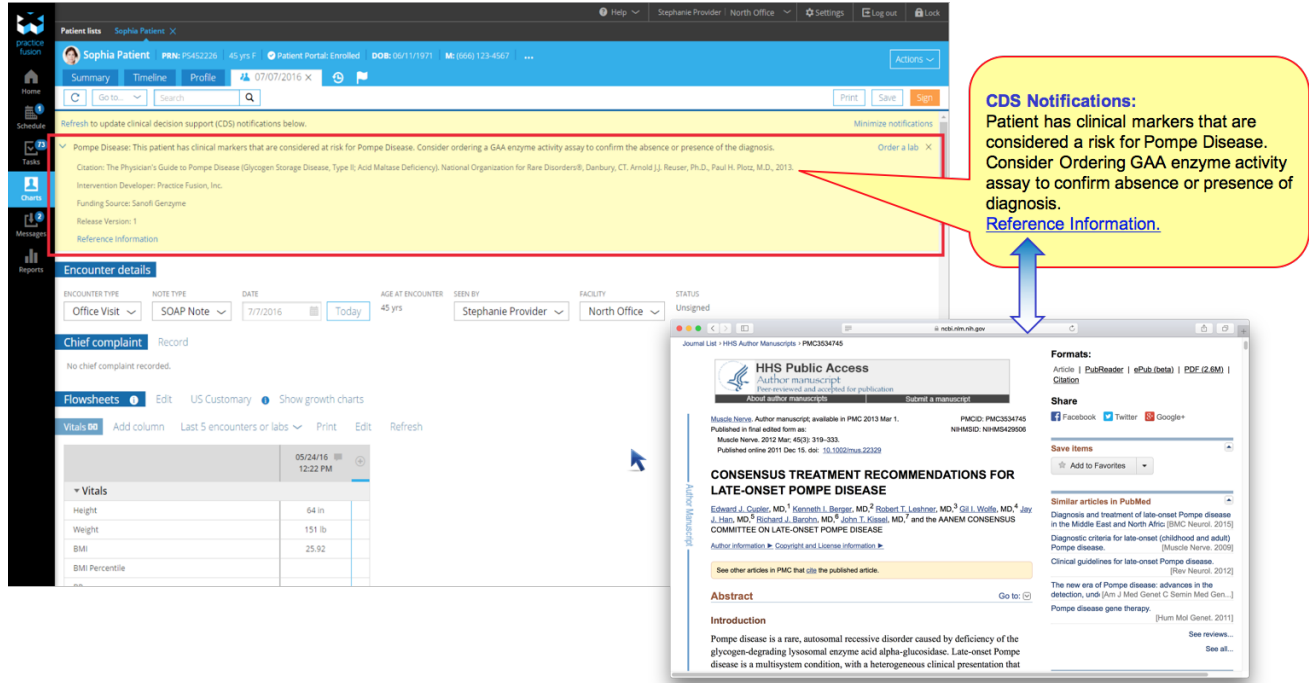


Figure 5-2: User Interface depicting CDS intervention which links to knowledge resource

Note that the contents of this alert have been magnified for the purpose of this illustration. Characteristics of this alert include:

- It appears at the top of the screen using fonts and colors that help to distinguish it from other content. (Alerts, by design, are intended to be noticed.)
- It includes mechanisms to process the intervention as appropriate (e.g. to acknowledge, accept or discard the alert). In this case, the alert may be closed (by clicking the X) or the suggestion to order a lab test may be accepted (by clicking on "order a lab").
- It provides a link to applicable reference information (as illustrated above by the PubMed^[3] screenshot.)
- It includes an option to "minimize notifications". This option allows the user to minimize the number of alerts displayed, by selecting the types of alerts they wish to receive in their user preferences.

Alert Fatigue

Alert fatigue^[4] is an unwanted side effect of CDS. Alert fatigue occurs when clinicians become overwhelmed by or desensitized to CDS alerts because of their sheer number, intrusive nature, or non-relevance to a clinical situation. The danger of alert fatigue is that the clinician will miss something important as a result. Strategies are required to minimize alert fatigue. Some of the interesting ideas proposed by thought leaders in CDS include:

- Increasing the specificity of alerts;
- Allowing users to customize CDS alerts by types of interventions
- Using a human factors approach to designing alerts

SNOMED CT is able to help with the first two items above. Firstly, it can be used to increase the specificity of the CDS conditions that trigger the alerts. And secondly, it can be used to distinguish between different types of interventions to enable customization to occur. Please refer to False Positives and False Negatives in the section Context in CDS Rules for more information on minimizing alert fatigue.

^[1] <https://www.ncbi.nlm.nih.gov/pubmed/>

^[2] Screenshot provided by Practice Fusion.

^[3] <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3534745/>

^[4] Uses content from: <https://psnet.ahrq.gov/primers/primer/28/alert-fatigue>

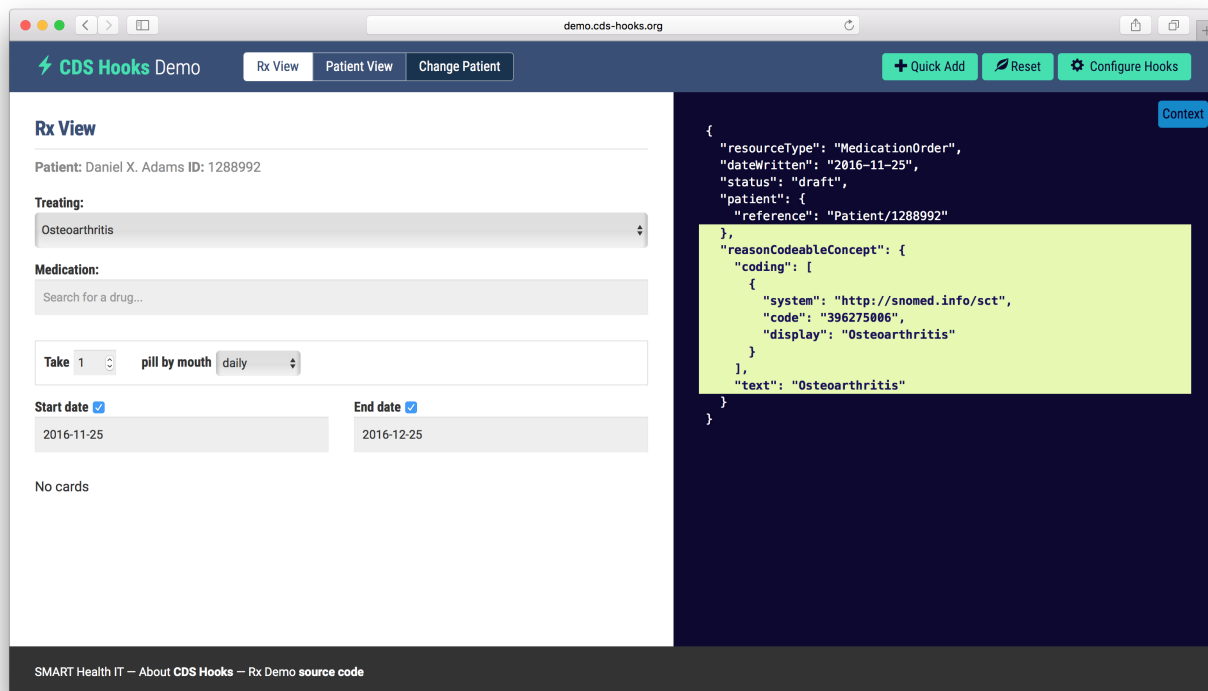
5.1. Standards for CDS Communications

An example of a standard which relates to CDS communications is provided below.

CDS Hooks

CDS Hooks is a relatively new CDS initiative which aims to automate the launching of applications that assist with decision support. CDS Hooks are designed around the premise of a clinician initiating a triggering activity within the EHR. When the triggering activity occurs, the EHR automatically sends a notification in real time to a decision support service (DSS).¹ This notification is considered the "hook" to the decision support logic. An example of a triggering activity would be a clinician writing a prescription. Some pre-defined hooks have already been developed and new hooks can be defined and added to the catalogue as required. Once the DSS is aware of the specific event, it may generate a response in the form of a "card" to be displayed in the UI of the EHR. An example of an "information" card might be one that contains pricing data about a proposed drug. The DSS could then propose a more cost effective "suggestion" card as an alternative. The other type of card the DSS may offer is an "app link" card which, as the names suggests, provides a link to an external application that can assist with further decision support. This architecture eliminates the need for the user to be aware of specific decision support applications that may be useful. The final outcome or choice, initiated by the app link card process, can then be automatically transferred to the appropriate field(s) in the EHR. A clinician has the option to accept or decline any suggestions present in the card. References to external knowledge resources may also be present in CDS Hooks cards.

The screen shot² below captures part of a CDS hook. Note that the condition the clinician is treating is represented using the SNOMED CT code for 396275006 |Osteoarthritis| .



The screenshot shows a web application interface for CDS Hooks. The main content area is titled "Rx View" and displays patient information (Daniel X. Adams, ID: 1288992) and a prescription for Osteoarthritis. The prescription details include a medication search field, dosage (1 pill by mouth daily), and dates (Start date: 2016-11-25, End date: 2016-12-25). A "Context" panel on the right displays a JSON object representing the CDS hook data:

```

{
  "resourceType": "MedicationOrder",
  "dateWritten": "2016-11-25",
  "status": "draft",
  "patient": {
    "reference": "Patient/1288992"
  },
  "reasonCodeableConcept": {
    "coding": [
      {
        "system": "http://snomed.info/sct",
        "code": "396275006",
        "display": "Osteoarthritis"
      }
    ],
    "text": "Osteoarthritis"
  }
}

```

Figure 5.1-1: CDS Hooks demonstration tool

For more information about CDS Hooks, please refer to <http://cds-hooks.org/>.

¹ HL7 definition of DSS: "A Decision Support Service takes in patient data as the input and provides back patient-specific assessments and recommendations."

² Screen shot from CDS Hooks Demo