

Chapter 6 / Capítulo 6

Primum non nocere: Software engineering
for medicine
ISBN: 978-9915-9876-8-2
DOI: 10.62486/978-9915-9876-8-2.Ch06
Pages: 95-110



© 2026 The authors. This is an open access
book chapter distributed under the terms of
the Creative Commons Attribution (CC BY)
4.0 License.

Chapter 6. Clinical Terminologies and Languages

Annier Jesús Fajardo Quesada  

Eduardo Antonio Hernández González 

René Herrero Pacheco 

Rolando Javier Álvarez Pérez 

1 Granma University of Medical Sciences. Granma, Cuba.

2 University of Medical Sciences of Pinar del Río. Pinar del Río, Cuba.

“Data without context is just a number; data with the wrong meaning *are* worse than no data at all. When a laboratory sends a result and the system interprets it as something it is not, when one medical record mentions ‘fever’ and another ‘pyrexia,’ when an algorithm interprets ‘heart attack’ where the doctor wrote ‘suspected heart attack,’ communication breaks down, and with it the possibility of providing safe and coordinated care.”

Let’s imagine for a moment that two doctors from different countries are trying to discuss a patient. One speaks Spanish, and the other speaks Mandarin. Neither understands the other’s language, but both need to share critical information to save a life. The situation is desperate and almost comical if it is not so serious. Now, let’s imagine that this same situation occurs daily, not between people but between computer systems. A laboratory sends a result labeled “glucose” to an electronic health record system that expects “GLU.” A pharmacy records a prescription for “paracetamol” that the dispensing system interprets as “acetaminophen.” An artificial intelligence algorithm trained on data from a Boston hospital receives data from a Madrid hospital and does not understand a thing because the coding systems are different. This is not a hypothetical scenario; it is the everyday reality of healthcare interoperability, and behind each of these miscommunications are patients receiving subpar care, doctors wasting time, and missed opportunities for improvement.

The solution to this problem is not technical but semantic. It is not enough for systems to be able to exchange data; they must exchange meaning. They must speak the same language or at least have a common dictionary that allows them to translate from one to the other. This is exactly what clinical terminologies are about: bringing order to the chaos of language, ensuring that when one system says “hypertension,” all the others understand exactly the same thing. It may seem like a minor detail, a matter of vocabulary of little importance, but it is actually one of the pillars upon which digital health is built. Without common terminology, there is no real interoperability; without interoperability, there is no integrated healthcare system; and without integration, technology fails to deliver on its promise of improving care.

In this chapter, we explore the fascinating world of clinical terminologies. We will see why it is so important that data mean the same thing across all systems, what problems arise when it does not, and how the major international terminologies address these challenges. The goal is not to turn the reader into an expert coder but to provide the necessary awareness to understand why this aspect—often overlooked in development projects—is actually critical for safety and quality. Because medical software that does not speak the same language as the rest of the system is not a solution, it is an additional problem.

6.1 WHY DATA MUST MEAN THE SAME THING ON BOTH SIDES

There is a natural tendency, especially among those who do not work directly with clinical information, to believe that data are objective and universal, that a number is simply a number and a word is a word and that their meaning is self-evident without the need for interpretation. This naive belief, which is already problematic in everyday life, becomes directly dangerous to patient safety in the healthcare setting. Clinical data are not abstract entities with universal and immutable meaning but rather complex human constructs laden with context, cultural nuances, implicit assumptions, and coding decisions.

When an experienced physician writes “fever” in an electronic health record, they are not simply recording an objective, neutral fact but actively interpreting a multifaceted clinical reality: they have previously measured a specific body temperature with a thermometer, considered the patient’s full context—including their age and baseline conditions—decided that this elevated temperature warrants being classified as “fever” rather than “hyperthermia” or simply “elevated temperature,” and they choose that specific term among other possible alternatives. When a different doctor reads that same record days later, they make a similar but not identical interpretation on the basis of their own clinical experience, their particular medical training, and the context of reading. This process of encoding and decoding, which works reasonably well among humans owing to shared medical education, a common clinical context, and the ability to ask for clarification, breaks down critically when automated computer systems intervene.

For a computer system lacking true semantic understanding, “fever” is merely a string of ASCII characters with no inherent meaning. It does not automatically associate a numerical range of body temperatures; it does not conceptually distinguish between a mild fever of 37.5°C and a high fever of 40°C; it does not know that it may be due to a bacterial infection or environmental

heatstroke; it cannot infer the therapeutic implications of that condition. If a different system receives that same text “fever” via data exchange, it stores it as-is in its database but cannot do anything intelligent or automated with that information: it cannot add it to comparable epidemiological statistics, it cannot generate alerts on the basis of temperature thresholds, and it cannot automatically correlate it with other clinical data from the patient to detect patterns.

For computer systems to truly do something useful with clinical information beyond simple storage, the data must be encoded in a standardized manner using formal medical terminologies that assign precise, structured, and shared meanings to each clinical term.

6.1.1 THE FIVE FUNDAMENTAL REASONS FOR SEMANTIC INTEROPERABILITY

The first fundamental reason why data must mean exactly the same thing across all systems that exchange information is direct patient safety. A classic and tragically common example is documented drug allergies. If a primary care system records that the patient is allergic to “sulfamides” using that colloquial term and another hospital system, upon receiving that information via electronic exchange, interprets it as the specific chemical group “sulfonamides”, or vice versa, the receiving system may fail to recognize through automatic matching that both terms conceptually refer to the same allergy and, therefore, may not generate the corresponding critical alert when a physician inadvertently prescribes a medication from that group. The result can be a severe, potentially anaphylactic allergic reaction that could have been completely avoided with standardized coding that eliminates terminological ambiguity.

Real patients do not receive medical care at a single, isolated healthcare facility but rather necessarily move between primary care, specialist clinics, emergency departments, acute hospitalization, posthospital rehabilitation, and often between different autonomous communities or even countries with distinct healthcare systems. For relevant clinical information to follow them reliably—such that the night shift emergency physician knows what treatments they have recently received and so that a specialist in a different city is aware of the results of previous diagnostic tests performed at another facility—computer systems must be able to “understand” one another semantically.

When population-based epidemiological studies are being conducted, when it is necessary to know precisely how many patients have a specific disease in a region or when the comparative effectiveness of different treatments is being evaluated through analyses of large databases, clinical data must be semantically comparable and aggregatable. If every hospital or system codes diagnoses in its own particular way, using its own nonstandardized local abbreviations and varying levels of granularity in the description, the results of any analysis will be inconsistent and biased, and the conclusions will potentially be erroneous or dangerous. Large healthcare databases that promise to revolutionize research through big data analysis are useless or even counterproductive if the data they contain do not speak the same precise semantic language. It is like trying to perform rigorous statistics by mixing responses in different natural languages without a translator or a common dictionary.

Doctors currently spend a large amount of their working time entering data into computer systems, and an equally considerable amount of interpreting data are entered by other colleagues who use different formats or terminologies. If the systems spoke the same standardized semantic language, much of that double-entry and reinterpretation work could be automated or eliminated. A laboratory result correctly coded with standard terminology can automatically populate timeseries charts, generate smart alerts when it falls outside safe ranges, and be incorporated into discharge reports without the need for manual rewriting. A precisely coded medication prescription can be automatically checked against known drug interactions, maximum recommended doses, and specific contraindications for that patient. However, for this intelligent automation to be possible, the data must be semantically interpretable by machines, not merely readable by humans who spend time decoding it.

AI models are primarily trained on historical data, and the semantic quality of that data directly determines the quality, reliability, and safety of the resulting models. If the training data are inconsistent, with ambiguous labels and variable coding that conceptually represents different things, the algorithms will learn erroneous patterns, artifacts, or—worse still—systematic biases that do not correspond to clinical reality. An algorithm trained to predict early sepsis that does not semantically distinguish between “fever” as a subjectively reported symptom and “elevated body temperature” measured objectively is useless or dangerous. An algorithm that learns from data where “myocardial infarction” means different things in different hospitals depending on their

local coding criteria will be unpredictable in its behavior. Rigorous semantic standardization is an absolute prerequisite for reliable and ethically acceptable medical artificial intelligence.

6.1.2 EXAMPLES OF AMBIGUITY AND TERMINOLOGIES AS A SOLUTION

To understand the practical magnitude of the problem of lack of standardization, it is helpful to examine real and common examples of terminological ambiguity in everyday clinical practice. A single disease can have multiple colloquial and formal names: “hypertension,” “HTN,” “high blood pressure,” “high blood pressure,” and “systemic hypertension.” A single drug may have different names depending on the country or tradition: “paracetamol” in Europe, “acetaminophen” in the Americas, and “APAP” in some technical contexts. The same surgical procedure can be described with different levels of clinical detail: “heart surgery” (very generic), “coronary bypass” (more specific), and “coronary artery myocardial revascularization” (technically precise). When computer systems handle these natural variations in medical language without a formal mechanism to semantically unify them, information becomes irreparably fragmented and loses its added value for coordinated patient care.

The practical consequences of this lack of semantic standardization are evident daily in hospitals and healthcare centers around the world. Physicians who lack important test results because the originating laboratory used a code different from what the receiving system expected. Nurses who administer the wrong medication because the electronic prescription used a brand name they did not recognize or that the system did not correctly link to the active ingredient. Researchers who spend months on tedious work “cleaning” data and manually correcting inconsistencies before they can actually begin an epidemiological study. Patients who have to repeat invasive diagnostic tests because the results from other centers cannot be automatically interpreted by the current system. All this massive waste of professional time, healthcare resources, and opportunities for improvement has a single, identifiable root cause: we have not yet managed to obtain healthcare IT systems to speak the same precise semantic language.

The solution to this problem necessarily involves adopting standardized medical terminologies, rigorously developed by international organizations with the participation of clinical experts from around the world. These formal terminologies are not simple flat lists of approved terms but complex ontological structures that organize medical knowledge systematically, with logical hierarchical relationships (hypernymy and hyponymy), documented synonyms, precise definitions of each concept, and unique, unambiguous codes for each clinical entity. When a computer system uses these standard terminologies, each piece of clinical data is associated not only with the readable text that the physician sees on the screen but also with a formal code that machines can interpret unambiguously in any other system that shares the same terminology. This code allows data to travel between heterogeneous systems while preserving its essential clinical meaning intact.

For medical software developers, understanding the critical importance of standardized terminologies is essential from the initial architectural design phase. This is not a minor technical detail that can be postponed for later, but a fundamental decision that determines the product’s future interoperability, its ability to integrate with other systems in the healthcare ecosystem, and the long-term utility of the generated data for research and quality improvement. Choosing one standard terminology over another—or worse, ignoring this aspect and using only unstructured free text—determines whether the product will be an isolated island of information or a connected node in the healthcare network. Deliberately ignoring this aspect means building a product that, no matter how technically sound it may be in its isolated functions, will exist disconnected from the broader healthcare ecosystem—like an island unable to communicate efficiently with the mainland.

The massive transition to semantically standardized systems is neither easy nor quick. It requires significant financial investment, extensive training of clinical and technical staff, profound changes to established workflows, and institutional coordination that often extends beyond the boundaries of a single organization. However, it is an investment with a guaranteed return in multiple dimensions. Every time clinical data can be reused without needing to be manually rewritten, every time a safety alert works correctly because the system semantically understood the documented allergy, every time a research study can begin without months of prior data cleansing, part of that initial investment is recouped. Above all, every time a patient receives safer, more coordinated care based on complete information because their data travel with them while retaining their meaning, the ultimate purpose of this entire collective effort is fulfilled: to truly

improve the health of the people who trust the healthcare system.

6.2 SNOMED-CT: THE SEMANTIC BACKBONE

Let us imagine that we want to build a universal library containing all of humanity's accumulated medical knowledge. Not just any traditional library but also one where every book, every article, and every piece of information is cataloged with pinpoint precision; where every medical concept has a specific place within an organized system; and where the logical relationships between ideas are clearly defined and can be navigated automatically. A library in which “pneumonia” was searched not only revealed all the documents mentioning that disease but also automatically displayed its hierarchical relationship to “respiratory infection,” its associated clinical manifestations such as “fever” and “productive cough,” its standard treatments such as “antibiotics,” and any other relevant medical concept semantically connected. This ideal library of medical knowledge exists not as a physical building with bookshelves but as a computable clinical terminology called SNOMED-CT, and it is likely the most ambitious, comprehensive, and sophisticated tool ever created by humanity to organize medical knowledge in such a way that machines can process it intelligently and humans can benefit from that organization.

SNOMED-CT, which stands for Systematized Nomenclature of Medicine Clinical Terms, is a comprehensive, multilingual, open-domain clinical terminology that covers most of the semantic knowledge necessary for modern healthcare practice.⁽¹⁾ Developed and rigorously maintained by SNOMED International, a nonprofit organization that brings together dozens of member countries and thousands of clinical experts, this terminology currently contains more than 350,000 unique and active concepts, organized into logical hierarchies and linked to one another through dozens of types of formal semantic relationships.

It is not simply an extensive list of approved medical terms, as a traditional dictionary might be, but a complex ontological semantic network that faithfully reflects the deep structure of medical knowledge: how diseases relate to their symptoms and signs, diagnostic and therapeutic procedures to their indications and contraindications, drugs with their mechanisms of action and adverse effects, and laboratory findings with the pathological conditions they suggest.

6.2.1 UNIQUE CONCEPTS, MULTIPLE DESCRIPTIONS, AND HIERARCHICAL STRUCTURE

What makes SNOMED-CT truly special is not only its impressive size and coverage but also its fundamental conceptual approach that clearly distinguishes between meaning and its linguistic expression. Each entry in SNOMED-CT is not merely a term or a word but also an abstract medical concept with a unique and established clinical meaning. That abstract concept is assigned a unique 18-digit numerical code that identifies it unambiguously and uniquely worldwide, regardless of language or context. However, the same concept may have multiple associated linguistic descriptions that express it in different ways: the preferred term in British English, the preferred term in American English, valid alternative synonyms, common clinical abbreviations, and official translations into dozens of national languages.

Thus, when a computer system sends a SNOMED-CT code, it actually sends a precise clinical meaning, not simply a word in a particular language. The receiving system can display the same meaning to the clinician in their preferred local language, using the term most familiar in that cultural context, but the key point is that both systems know with certainty that they are referring to exactly the same clinical entity, not ambiguous interpretations.

The hierarchical structure of SNOMED-CT allows for systematic navigation from the most general to the most specific, or vice versa, in a logical and computable manner. All the concepts are organized into 19 main hierarchies on the basis of “is-a” taxonomic relationships. For example, the specific concept “bacterial pneumonia caused by *Streptococcus pneumoniae*” is a more specific type of “bacterial pneumonia,” which in turn is a type of “pneumonia” (inflammation of the lung parenchyma), which is a type of “infectious disease of the lower respiratory tract,” which is a type of “infectious disease,” which is ultimately a type of “disease” or pathological condition.

This formal hierarchical structure allows computer systems to automatically make useful logical inferences: if a patient has a specific bacterial pneumonia documented, the system can automatically deduce that they have pneumonia in general, that they have an infectious disease and that they have a disease requiring treatment. This is extraordinarily useful for information retrieval in searches (“find all patients with any type of respiratory infection”), for epidemiological

data aggregation (“how many cases of infectious disease did we have this month”), and for generating rule-based clinical alerts (“alert if the patient has an infection and has not been prescribed antibiotic therapy”).

However, SNOMED-CT goes far beyond simple taxonomic classification hierarchies. It also includes other types of semantic relationships that connect concepts from different hierarchies to one another, creating a rich and multidimensional network of medical knowledge. For example, the concept “pneumonia” can be semantically related to the concept “fever” via a formal “has associated clinical finding” relationship and to the concept “amoxicillin” via a “has drug treatment of choice” relationship. These attributive relationships allow for the representation of causal and associative medical knowledge in a much richer and more useful way than a simple classification taxonomy.

A system that uses SNOMED-CT in a sophisticated manner can answer complex clinical questions that span multiple domains: “Find all patients with an infectious disease of the respiratory tract presenting with high fever who have been treated with beta-lactam antibiotics in the past 30 days.” Without a terminology featuring these formal semantic relationships, answering that question would require imprecise text searches in free-text notes or multiple separate queries to noninteroperable databases.

6.2.2 TECHNICAL INTEGRATION, CONTINUOUS UPDATES, AND DATA QUALITY

For medical software developers, SNOMED-CT offers an invaluable resource that avoids reinventing the semantic wheel. Instead of having to create and maintain their own lists of medical terms—with all the clinical consensus work and potential for systematic error that entails—they can use this terminology, which has already been developed, validated, and is continuously maintained by international experts across multiple specialties. This not only saves months or years of development work but also ensures that their products will be semantically interoperable with other systems that also use SNOMED-CT correctly. A code sent from their application will be understood unambiguously by any other system that adheres to the standard, regardless of the country of origin, the local implementation language, or the specific software vendor.

The technical integration of SNOMED-CT into a medical software application is not trivial given the complexity of the terminology, but there are tools and standards that considerably facilitate the task. There are open-source and commercial software libraries available for virtually all modern programming languages (Java, Python, C#, and JavaScript) that allow you to query the terminology locally, validate entered codes, programmatically navigate hierarchies, and retrieve descriptions in different languages on the basis of the user’s context. There are also specialized terminology servers that expose standardized APIs (such as FHIR Terminology Services) so that applications can query SNOMED-CT remotely via the web without needing to incorporate the entire knowledge base into their local installation. The choice between one architectural option or another depends on the specific needs of each project: applications that must operate in environments without an internet connection (such as in rural areas or during emergency deployments) may require a complete local copy of the terminology; applications that are always connected can benefit from centralized services that update automatically without user intervention.

For physicians and clinical professionals involved in designing applications that use SNOMED-CT, understanding the basic characteristics of this terminology helps them grasp why computer systems sometimes offer coding options that seem limited or why they do not immediately find terms they consider obvious and common. Formal terminology is not a perfect reflection of colloquial medical natural language but rather a standardized artificial construct with its own logical rules, naming conventions, and coverage limitations. It may be that the specific term the clinician routinely uses in their practice is not exactly the preferred term in SNOMED-CT or that the clinical concept they wish to express precisely does not have an exact entry in the terminology and they must settle for a more general, approximate concept. Understanding these reasonable limitations helps in having realistic expectations about what the technology can currently offer and in collaborating constructively with developers to find the best available coding solution.

An important point for developers to understand is that SNOMED-CT is neither static nor immutable. The terminology is continuously updated through biannual cycles to incorporate new emerging medical knowledge, correct detected errors, improve the coverage of underrepresented specialties, and refine existing semantic relationships. A new official international version is published every six months, and national member countries may also create local extensions to

address specific needs of their particular healthcare system (such as specific funding procedures or locally available medications). This means that developers must stay up to date with these updates, implementing system update processes to use the latest versions when appropriate and ensuring backward compatibility with previous versions when required by regulatory requirements.

The relationship between SNOMED-CT and other specialized medical terminologies is also relevant for realistic implementation. SNOMED-CT does not aim to be the only term used in the healthcare ecosystem, nor does it seek to replace all others; rather, it is intended to serve as the central semantic backbone that connects other, more domain-specific terminologies. For example, for medications, another term called RxNorm (in the U.S.) or ATC (internationally) is typically preferred; for clinical laboratory tests, LOINC is the dominant standard; for surgical procedures in billing contexts, CPT or the ICD-10-PCS are preferred.⁽²⁾ SNOMED-CT can conceptually act as a reference map that relates and translates between all of them through formal mappings, allowing a concept originally expressed in a specific terminology to be automatically translated into the others depending on the context of use. This capacity for terminological integration is essential in a real healthcare ecosystem where multiple historical and specialized standards necessarily coexist.

For development teams, an important architectural decision is determining what level of granularity and specificity to use from SNOMED-CT for different use cases. The terminology allows coding with varying degrees of clinical specificity, ranging from very general to extremely detailed and specific. For example, one can simply code “diabetes mellitus” (general), or one can specify precisely “type 2 diabetes mellitus with stage 3 renal complications” (very specific). The choice of the appropriate level of detail depends critically on the subsequent use of the clinical data. For an electronic health record that must support precise clinical decisions and longitudinal follow-up, specific details are important and valuable. For an administrative billing system based on diagnosis-related groups, the general category may suffice.

A terminology, no matter how comprehensive and sophisticated, does not save lives on its own. What truly saves lives is the ability to share accurate information among professionals, to make clinical decisions informed by complete data and to coordinate care seamlessly across multiple professionals and levels of care. SNOMED-CT is the semantic infrastructure that makes all of this possible in today’s digital world, but the true value lies in what developers, clinicians, and organizations build upon it. Developers who understand this in depth do not view the terminology as an imposed technical burden but rather as the fundamental semantic infrastructure upon which they can build innovative applications that truly make a difference in patients’ lives.

6.3 LOINC FOR LABORATORY AND RXNORM/ATC FOR MEDICATIONS

Imagine a patient who is receiving a blood test at a clinical laboratory in their city. The results revealed a glucose concentration of 110 mg/dL, which the doctor interpreted in context. That simple number travels from the automated analyzer to the laboratory management system, and from there, it must be transmitted to the electronic health record so that their primary care physician will later consult to make treatment decisions. However, for that number to be truly clinically useful, it is not enough for it to arrive correctly over the network; the receiving system must know with certainty that “110” refers specifically to blood glucose, not total cholesterol or creatinine; that it is expressed in mg/dL, not in mmol/L, which would require conversion; that it is a venous blood sample obtained by venipuncture, not a fingerstick capillary sample; and that it is a measurement taken under fasting conditions, not postprandially after a meal. All this critical contextual information, which is obvious to laboratory professionals because of their training, is completely inaccessible to the receiving computer system if it is not properly encoded using agreed-upon standards. This is where two specialized terminologies come into play, together with SNOMED-CT as the backbone, to complete the fundamental semantic ecosystem of modern digital health: LOINC for clinical laboratory tests and RxNorm or ATC for medications.

LOINC, which stands for Logical Observation Identifiers Names and Codes, is the international reference terminology for unambiguously identifying laboratory tests and other structured clinical observations.⁽³⁾ Developed and rigorously maintained by the Regenstrief Institute at Indiana University, LOINC provides a unique alphanumeric code for each type of test or observation, with sufficient semantic granularity to distinguish all variants that may be clinically relevant in real-world practice. Blood glucose is not the same as 24-hour urine glucose, nor is a fasting measurement the same as a postprandial oral glucose tolerance test, nor is a serum

measurement the same as an EDTA plasma measurement or a capillary blood measurement. LOINC allows for the systematic coding of all these contextually relevant variants in a precise and standardized manner, ensuring that when a laboratory sends a result coded with LOINC, the recipient can interpret it correctly without ambiguity regarding exactly what was measured and under what preanalytical conditions.

6.3.1 THE STRUCTURE OF LOINC AND ITS PRACTICAL IMPLEMENTATION

The logical structure of the LOINC is based on six main axes that, when systematically combined, uniquely and unambiguously define each clinical or laboratory observation. The component specifies which substance or analyte is measured (such as glucose, creatinine, or hemoglobin). The property indicates which aspect of the component is quantified (mass concentration, molar concentration, cell count, partial pressure). The time frame specifies whether the measurement is a single point in time or integrated over a time interval. The system or sample type identifies the biological material (venous blood, serum, citrated plasma, urine, and cerebrospinal fluid). The analytical method is specified when clinically relevant for interpretation (enzymatic method, immunoturbidimetry, chromatography). Finally, the scale indicates the type of result (numerical quantitative, ordinal qualitative, descriptive nominal, or free narrative). This systematic approach to semantic decomposition allows LOINC to currently contain more than 90,000 standardized terms,⁽³⁾ covering not only conventional laboratory tests in biochemistry and hematology but also structured clinical observations, coded vital signs, functional assessment scales, structured imaging results, and many other types of clinical data requiring precise identification.

For clinical laboratories, adopting LOINC as an internal coding standard is a fundamental step toward true semantic interoperability. When a laboratory sends results to other systems using standardized LOINC codes, any recipient that understands the standard will be able to interpret them correctly and automatically, regardless of how the sending laboratory internally refers to those tests in its own legacy system or commercial catalog. This is especially critical when patients move between different care centers throughout their lives or when data from multiple sources are aggregated for epidemiological research or population-based quality management. Without LOINC, each laboratory inevitably uses its own local internal codes, and semantic integration becomes a logistical nightmare of adhoc manual mappings that consumes valuable professional time, is prone to interpretation errors, and quickly becomes outdated when catalogs change.

For software developers working with laboratory data in clinical applications, LOINC offers a powerful way to standardize information from multiple heterogeneous sources. A medical record application that receives results from dozens of different laboratories can, if all correctly use LOINC, process them in a completely uniform manner internally, applying the same clinical interpretation rules, the same population-adjusted normal reference ranges, and the same critical value alerts. Without this prior terminological standardization, the application would have to know and keep updated the proprietary internal codes of each connected laboratory and apply potentially different interpretation rules for each one—which is technically possible but extremely complex, costly to maintain, and prone to inconsistencies that compromise patient safety.

The implementation of LOINC in a healthcare organization involves not only a technical code-mapping project but also a disciplined exercise in standardizing clinical processes and quality management. It compels laboratories to precisely and explicitly define which tests they actually perform and how they name them internally, to resolve accumulated historical ambiguities and to align with internationally agreed-upon practices. This standardization process, although costly in terms of time and resources initially, has substantial and long-term benefits in terms of data quality, interanalyst consistency, and operational efficiency. A laboratory that knows exactly which tests it offers, uses which methodology, and how it is named in the international standard, is better positioned to rationally manage its service catalog, adequately train its technical staff, and communicate transparently with prescribing clinicians.

6.3.2 RXNORM: DRUG STANDARDIZATION WITH FLEXIBLE GRANULARITY

Let us now turn to the domain of medications, which is, if anything, even more semantically complex than the laboratory domain because of the multiplicity of dimensions involved. A single drug may have multiple commercial formulations with different dosages, multiple active

ingredients in fixed combinations, multiple trade names protected by patents in different countries, and multiple routes of administration and dosage forms. To address this inherent complexity, there are two main reference terminologies, with complementary approaches that address different usage needs: RxNorm, developed and maintained by the U.S. National Library of Medicine (NLM), and the ATC (Anatomical Therapeutic Chemical) classification, maintained by the World Health Organization (WHO).⁽⁴⁾

RxNorm is a terminology system that standardizes the names of medications in the U.S. market, providing stable, unique identifiers for each relevant pharmacological concept. What makes RxNorm particularly powerful is that it explicitly distinguishes between different levels of semantic abstraction required for various clinical uses: from the pure active ingredient (such as “atorvastatin” as a chemical substance) through combinations of active ingredients (such as “atorvastatin + amlodipine” as a conceptual product) to specific commercial formulations with defined dosages and dosage forms (such as “atorvastatin 20 mg coated tablets”). Each of these levels of abstraction is assigned its own unique identifier (RXCUI), and RxNorm formally maintains the semantic relationships between them so that one can systematically navigate from the generic to the specific and vice versa according to the needs of the use case.

This multilevel granularity structure is extremely useful for various clinical applications. An electronic prescribing system can operate comfortably at the generic active ingredient level, allowing the physician to prescribe “atorvastatin” without having to specify the brand name or exact formulation, leaving that decision to the dispensing pharmacist. A hospital pharmacy dispensing system may require a specific formulation level to know exactly which physical product to retrieve from the warehouse. A drug interaction alert system can function efficiently at the ingredient level, regardless of the specific commercial formulation the patient is taking. RxNorm allows all these different systems to speak the same semantic language, each operating at the level of detail it needs for its specific function, while maintaining consistency among them.

6.3.3 ATC: THERAPEUTIC CLASSIFICATION FOR POPULATION ANALYSIS

ATC classification, on the other hand, has a different and complementary conceptual approach. It does not focus primarily on drug names or their commercial presentations but rather on their therapeutic classification on the basis of the organ or biological system on which they primarily act and their specific pharmacological properties. The ATC structure () assigns each active ingredient a five-level alphanumeric code on the basis of its primary therapeutic indication. For example, all drugs for cardiovascular diseases begin with the letter C; within this group, C10 specifically refers to lipid-lowering agents; C10AA refers to statins as a pharmacological group; and C10AA05 specifically refers to atorvastatin as an individual active ingredient.

ATC is primarily used for epidemiological studies of drug utilization, analyses of population-based prescribing patterns, pharmaceutical formulary management, and postmarketing safety surveillance. It allows drugs to be grouped by functional therapeutic categories, facilitating aggregate analyses that would not be possible using the individual names of thousands of different commercial formulations. For example, to study total antibiotic consumption in a population over a season, ATC allows for the easy grouping of all drugs in group J01 (antibacterials for systemic use), regardless of their specific active ingredient or commercial formulation. For pharmacovigilance studies and drug utilization analyses at the public health level, this therapeutic grouping capability is absolutely essential.

For medical software developers, the choice between RxNorm and ATC (or the decision to use both simultaneously) depends critically on the application's specific use case and the geographic context of implementation. A clinical electronic prescribing application will likely need the detailed granularity of RxNorm (or its national equivalent) to handle formulations, dosages, and routes of administration accurately. A public health application focused on analyzing aggregated population data may prefer the therapeutic perspective of ATCs for trend analysis and benchmarking. In many cases, the ideal approach is to use both terminologies in a complementary manner, maintaining mapping relationships between the codes of one and the other in the master database. A specific medication has an RxNorm code that uniquely identifies it on the market, as well as an ATC code that classifies it therapeutically, and both pieces of information are useful in different application contexts.

The integration of these specialized terminologies with SNOMED-CT as the backbone completes the interoperable semantic ecosystem. SNOMED-CT provides a general reference framework, with concepts for diseases, symptoms, procedures, laboratory findings and medications

at a clinical abstraction level. LOINC and RxNorm/ATC provide the specific technical details within their respective specialty domains. The formal relationships between them are defined in official reference maps maintained by the standards organizations so that one can navigate semantically from one terminology system to another as needed. For example, a SNOMED-CT concept for “type 2 diabetes mellitus” may be formally related to the specific LOINC code for “glycated hemoglobin” as a recommended follow-up test and to the ATC codes corresponding to oral antidiabetics as first-line treatment. This rich network of semantic relationships between terminologies is the technical foundation upon which true semantic interoperability in the digital health ecosystem is built.

For software development teams, working effectively with these terminologies requires familiarity with their internal logical structures and with the technological tools available for their practical management. There are standard web APIs for querying LOINC and obtaining its full descriptions, RESTful web services for accessing RxNorm and its -level relationships, and open-source software libraries for handling ATC codes in local applications. There are also specialized semantic mapping services that help convert proprietary local codes from legacy systems to these international standards, facilitating gradual technological migration. The initial investment in learning to use these tools and integrating them architecturally pays off quickly in terms of the quality of the clinical data generated, ease of integration with other systems, and the capacity for subsequent population analysis.

A critical practical aspect of managing these terminologies is version control and the management of terminological changes. All these terminologies are periodically updated with new versions: LOINC has three releases per year (January, June, and September); RxNorm is updated monthly with new commercial products; and ATC is revised annually by the WHO. Systems that use them must have technical mechanisms and organizational processes to incorporate these updates without losing backward compatibility with historical data coded using previous versions. This is especially critical when longitudinal patient data spanning years are handled: a code that was valid and standard five years ago may have been refined, split, or even declared obsolete in more recent versions, and the system must be technically capable of interpreting it correctly within its historical context. Managing terminological change over time is a complex process that requires advanced architectural planning and appropriate versioning tools.

For physicians and clinical professionals who use applications based on these terminologies, the user experience should be completely transparent and should not increase their workload. They should not have to memorize or know the LOINC codes for the laboratory tests they routinely order or the ATC codes for the medications they prescribe daily. The system should provide them with intuitive search interfaces using common names in their language, with intelligent autocompletion, and handle the translation of those textual searches into the corresponding standard codes internally. The underlying terminological complexity must be carefully hidden beneath a layer of well-designed usability that respects natural clinical workflows. When this is effectively achieved, clinicians reap all the benefits of semantic standardization (interoperable data, smart alerts, quality analysis) without having to pay the price of learning a new artificial technical language.

The final quality of clinical data coded with these terminologies depends, as always in information systems, on those who enter it at the point of care. A physician who prescribes “paracetamol 500 mg tablets” is generating information that can be automatically coded by the system if it is well designed with synonym dictionaries. However, if the physician uses nonstandardized local trade names or enters dosages in inconsistent formats (sometimes “500 mg,” other times “0.5 g,” and still others “two tablets”), automatic coding becomes more difficult and prone to errors. That is why it is so important for systems to actively guide data entry through interfaces that suggest standardized options, validate semantic consistency in real time, and provide immediate feedback when they detect something inconsistent or uncodable.

Behind every LOINC code lies years of consensus-building work by laboratory experts from around the world to define exactly what a specific test means and how it differs from similar ones. All this international collaborative work—invisible to most clinical professionals who simply use the resulting systems—is the fundamental semantic infrastructure upon which all modern digital health is built. Without it, clinical data would be unintelligible noise; with it, it becomes actionable information for providing better and safer patient care. In addition, that, at the end of the day, is the only thing that truly matters in medicine.

6.4 MAPPINGS AND TRANSFORMATIONS BETWEEN TERMINOLOGIES

Imagine a chronic patient arriving at an emergency department in a city other than their own, far from their usual primary care physician. They bring with them a recent discharge summary from their home hospital, which documents that they have “type 2 diabetes mellitus with stage 3 renal complications” and are on active treatment with “metformin 850 mg tablets every 12 hours.” The emergency physician urgently needs to incorporate this critical information into the receiving hospital’s electronic health record to make safe management decisions, but this hospital’s computer system uses SNOMED-CT for diagnosis coding and ATC for medication coding, whereas the referring hospital uses ICD-10-CM for billing diagnoses and its own nonstandardized internal code for medication prescriptions. The patient’s information is clinically invaluable for continuity of care, but it is expressed in a terminological “language” that the receiving system cannot understand directly without translation. Someone—ideally a high-quality automated process—must semantically translate these concepts between systems. This is precisely what mappings between terminologies are all about: establishing formal semantic bridges that allow clinical information to travel from one system to another while preserving its essential meaning for decision-making.

A terminological mapping is, in essence, a declared correspondence between concepts in different medical coding systems.

It can be a simple one-to-one relationship, where a concept in one terminology exactly and unambiguously corresponds to a concept in another; a one-to-many relationship, where a general concept in one terminology corresponds to several more specific concepts in another, requiring a decision on which one to choose; a many-to-one relationship, where several specific concepts in the source terminology are necessarily grouped under a more general one in the target owing to a lack of granularity; and even more complex relationships may exist, with partial overlaps in meaning, approximate or “sufficiently similar” equivalences, or even a total lack of correspondence when a concept simply does not exist in the target terminology because it is too specific or novel.

Managing this inherent complexity is among the most fascinating and critical technical, semantic, and epistemological challenges of modern digital health engineering.

6.4.1 THE THREE REASONS FOR THE NEED FOR MAPPING

The practical need for terminological mappings arises from several fundamental reasons that coexist in the real-world healthcare ecosystem. The first is historical and legacy-related: different national and regional healthcare systems have adopted different standard terminologies over the decades on the basis of their legislative and technological contexts and now need to interoperate without replacing their entire legacy of historical data. A hospital may have been using ICD-9-CM for decades to code discharge diagnoses, and migrating its entire historical archive of millions of episodes to SNOMED-CT would be a titanic data engineering task with prohibitively high costs. Well-designed mappings allow new systems based on modern terminologies to understand and interpret old information coded with previous standards and enable historical systems to communicate bidirectionally with new systems without needing to replace the entire existing infrastructure all at once, facilitating gradual technological transitions.

The second reason is functional and relates to use cases: different medical terminologies inherently serve different clinical and administrative purposes, and a single information system may need to use several of them simultaneously for different functions. A single patient diagnosis may need to be coded in SNOMED-CT with high clinical granularity for detailed electronic health records, in ICD-10-CM with billing codes for the hospital discharge coding system, and in a specific simplified classification for a national epidemiological registry of a particular disease. Instead of forcing the clinician to manually enter the same diagnostic concept multiple times in different codes for different systems—which would be inefficient and prone to inconsistencies—automatic mappings allow systems to automatically translate the concept into the other representations needed for various downstream uses once it has been entered and validated in primary terminology.

The third reason is global interoperability and cross-border care: When clinical information is exchanged between countries with different terminological traditions or between systems that have adopted different national standards for historical reasons, mapping is the only technically viable way to ensure that data are mutually understood without critical semantic loss. A European

patient receiving emergency care in the United States, an international multicenter research study involving hospitals on different continents with different coding standards and an international drug safety alert requiring rapid and accurate data sharing among regulatory agencies in different countries—all these critical scenarios depend fundamentally on the existence of reliable, validated, and maintained mappings between the terminologies used in each jurisdiction.

The most important, reliable, and widely used mappings internationally are those officially developed and maintained by the organizations responsible for the source and target terminologies themselves. SNOMED International, for example, maintains and publishes updated official mappings between SNOMED-CT and the ICD-10 for both the detailed clinical version and the simplified mortality statistics version. The U.S. National Library of Medicine (NLM) maintains official mappings between RxNorm and various national drug terminologies used in other European and Asian countries. These official mappings are the result of meticulous, methodologically rigorous work by clinical experts and terminologists who analyze concepts by concept, determine the most appropriate semantic correspondence in each case, explicitly document ambiguous cases or instances of imperfect correspondence, and validate the results with real-world test data.

6.4.2 LOCAL MAPPINGS, GRANULARITY MANAGEMENT, AND DIRECTIONALITY

However, official international mappings, although of high quality, do not cover all the practical needs of real-world healthcare systems. Often, individual healthcare organizations need to create and maintain their own “local” mappings to connect their proprietary legacy systems with modern international standards or to adapt generic international standards to their specific local clinical practice needs. A hospital may have its own internal coding system for hundreds of laboratory tests, developed organically over decades of service evolution, and needs to systematically map it to standard LOINC to exchange results with other centers in the regional network. These local mappings are as important operationally as official international mappings are and require equally rigorous methodologies to ensure their semantic quality and ongoing maintenance.

Creating high-quality terminological mapping is not a trivial task that can be fully automated without expert supervision. It requires in-depth knowledge of both terminologies involved in their entirety, a contextual understanding of the specific clinical environment in which the mapped codes will be used, and a documented systematic methodology for resolving ambiguous or conflicting cases. Multidisciplinary teams that develop high-quality mappings typically follow a formalized process that includes a preliminary comparative analysis of the structures of both terminologies, automated identification of obvious correspondences on the basis of term similarity, manual review of doubtful or conflicting cases by clinical experts in the relevant specialty, empirical validation of the resulting mapping using real test data, and comprehensive documentation of the decisions made and cases of imperfect correspondence. It is a technically complex, semantically meticulous task that requires constant updating as the source and target terminologies evolve with new versions.

A crucial aspect of mapping design is the explicit management of differences in granularity between terminologies. When one term is systematically more detailed than another in a specific domain, one-to-many relationships may arise that require algorithmic or human decisions on how to proceed. For example, SNOMED-CT has very specific concepts for different types of diabetes mellitus and their specific complications, while an older classification, such as ICD-9, may have only “diabetes mellitus” as a general aggregate category. In these cases, of granularity asymmetry, a strategic decision must be made: either map everything to the general level, thereby losing valuable detailed information, or create contextual rules that allow for preserving detail whenever possible through dual coding. There is no single universal answer; it depends on the specific use case and the criticality of the detail for the clinical or administrative decision that will be supported by the data.

The directionality of the mapping also matters significantly and must be explicitly documented. A mapping that works well from Terminology A to Terminology B may not work equally well from B to A, especially when there are asymmetric relationships in terms of granularity or specialization. That is why good professional mappings are usually bidirectional when possible or at least clearly document the limitations and information loss in each translation direction. A system that receives data coded in a source terminology and needs to translate it into a target terminology must explicitly know whether the correspondence is exact (equivalence), approximate

(similarity), or whether there is simply no valid correspondence for that specific concept.

6.4.3 TECHNICAL IMPLEMENTATION, UNCERTAINTY MANAGEMENT, AND QUALITY

For medical software developers, working with terminology mappings in their applications involves architecturally incorporating the ability to perform robust automatic semantic translations. This can be implemented in various technical ways depending on the context: using remote terminological web services that offer updated official mappings (such as those from the NLM or SNOMED International), incorporating local mapping tables into system databases for quick lookups without network dependency, or implementing more complex transformation engines that apply algorithmic inference rules for cases not directly mapped. The architectural choice depends on the volume of data to be processed, the required frequency of real-time or batch translations, and the need to maintain automatic synchronization with periodic updates to the source and target terminologies.

A significant and often underestimated technical challenge is the explicit management of the uncertainty inherent in mappings. Not all terminological correspondences are perfect or exact; some are reasonable approximations, others are partial with a loss of nuance, and others simply do not exist for very specific or new concepts. Information systems that use mappings must be designed to handle this uncertainty transparently, for example, by flagging automatically translated data with metadata such as “automatically mapped—review recommended” so that the clinical user knows it may not be exact or by implementing mandatory human review workflows for cases flagged as questionable by the mapping algorithm. Explicit transparency regarding the quality and reliability of mappings is essential to maintain professional trust in the translated data and to enable the detection of translation errors before they cause harm.

The quality of mappings in production can be measured and monitored using objective metrics such as accuracy (the percentage of machine translations that are correct according to expert validation) and coverage (the percentage of source concepts that can be translated to the target without unacceptable loss of information). Teams maintaining operational mappings must continuously monitor these metrics, investigate errors detected through auditing, and iteratively improve their maps. It is a task of continuous maintenance throughout the system’s lifecycle, not a one-time implementation project with a defined end.

For physicians and clinical professionals who use systems based on terminology mapping, the ideal user experience should be completely seamless within their normal workflow. They should not have to consciously worry about whether the diagnosis they enter into SNOMED-CT is correctly translated to the ICD-10 for billing; the system should do this automatically and display both codes for verification if necessary. However, they should be aware, as part of their professional digital literacy, that such automatic translation exists and that it may have semantic limitations in complex cases. When a system displays clinical information from another healthcare facility, automatically translated from another terminology, they should know that they are viewing a semantic interpretation, not necessarily the exact original record from the clinician who treated the patient. This critical awareness helps you maintain an appropriately skeptical attitude and detect potential translation errors that could affect patient safety.

A particularly complex special case of mapping is the transformation between different versions of the same terminology over time. When SNOMED-CT releases a new version every six months (), some concepts may change codes because of refinement, disappear because of obsolescence, or split into several more specific concepts because of changes in the ontological model. Systems that handle longitudinal patient data accumulated over years require mappings between versions to correctly interpret older codes within the context of current terminology. These “medical record” mappings are provided by the organizations responsible for the terminologies, but their technical implementation requires special care to avoid the loss of valuable historical information or the introduction of anachronistic errors in retrospective interpretation.

The relationship between the quality of the mappings and the quality of the resulting clinical data is direct and critical. An incorrect mapping, especially if it is systematic and undetected, can turn clinical data that are perfectly valid in the source terminology into erroneous or misleading information in the target terminology, with potentially serious consequences for patient care decisions. That is why mappings must be treated with the same rigor in quality and security management as any other critical component of the medical information system. They must be developed by multidisciplinary teams of experts, empirically validated with real test data, subjected to formal version control, periodically audited by independent third parties, and

thoroughly documented. Clinicians' confidence in automatically translated data must be as high as their confidence in the original manually entered data, and this can be achieved only through rigorous quality processes.

As terminologies mature and become internationally standardized and as new systems predominantly adopt common standards such as SNOMED-CT and LOINC, the acute need for complex mappings may gradually diminish, but it will never disappear completely. There will always be legacy systems that must be kept operational for years; there will always be legitimate local needs for specific coding; and there will always be new emerging domains of medical knowledge (such as genomics or personalized medicine) that will require new specialized terminologies, which will then need to be mapped to existing standards. Mappings are, and will remain for the foreseeable future, an essential and permanently necessary part of the semantic infrastructure of global digital health. Software developers, data engineers, and clinicians who deeply understand their importance, learn to design them rigorously, and develop the skills to work with them critically are better prepared to build information systems that can truly communicate with one another securely, overcoming the barriers of medical language and allowing clinical information to flow seamlessly to where it is most needed: at the patient's bedside, at the moment of critical clinical decision-making.

REFERENCES

- What is SNOMED CT? · NLM Customer Support Center [Internet]. [cited 2026 May 13]. Available from: <https://support.nlm.nih.gov/kbArticle/?pn=KA-04795>
- Fung KW, Xu J, Brear H, Lane A, Lau M, Wong A, et al. Promoting interoperability between SNOMED CT and ICD-11: lessons learned from the pilot project mapping between SNOMED CT and the ICD-11 Foundation. *J Am Med Inform Assoc* [Internet]. August 1, 2024 [cited May 13, 2026];31(8):1631. Available from: <https://pmc.ncbi.nlm.nih.gov/articles/PMC11258399/>
- Health Data Standards and Terminologies: A Tutorial [Internet]. [cited 2026-05-13]. Available from: <https://www.nlm.nih.gov/oet/ed/healthdatastandards/02-530.html>
- RxNorm Overview [Internet]. [cited May 13, 2026]. Available from: <https://www.nlm.nih.gov/research/umls/rxnorm/overview.html>