Chapter 7
Clinical Terminology

Abstract Unlike most sciences, medical terminology is poorly structured. This creates major problems for semantic interoperability, where terms need to be used in a precise and unambiguous way. This chapter introduces the core concepts of clinical terminology, sets out a list of requirements (desiderata) and illustrates these with the story of the Chocolate Teapot.

Keywords Terminology • Synonym • Homonym • Code • Classification • Hierarchy • Concept • Coding scheme • Display term • Relationship • Value set • Identifier • Reference terminology • Interface terminology • Ontology • Expression • Vocabulary • Polyhierarchy • NEC (not elsewhere classified) • NOS (not otherwise specified) • Redundancy • Chocolate teapot

Why Clinical Terminology is Important

When, in the fifteenth century, Gutenberg's invention of the movable type led to the mass production and dissemination of books and written information, language was still relatively unformalised. It took until the eighteenth century before the great dictionaries and nomenclatures such as Dr Johnson's English Dictionary and Linnaeus' biological taxonomy were produced.

Sciences such as biology and chemistry have an internationally agreed formal structure for their terminology. Every living organism has a generic and specific Latin name expressed within a comprehensive biological taxonomy, which in many ways anticipated the full understanding of the evolution of life. All chemical structures are expressed in internationally standardised ways.

Medical terminology escaped formalization, leading to problems of ambiguity that are now recognised as a significant risk to patient safety. The lack of agreed medical terminology has been recognised as an issue for at least 250 years. There is even an old word, "nosology", to describe the development of medical terminology, but the need has increased with the use of computers. Modern nosologists call themselves clinical terminologists.

The historical, eclectic and ad hoc origins of medical terminology have encumbered anyone interested in health-care with the need to learn a whole new language,

T. Benson, G. Grieve, *Principles of Health Interoperability*,
Health Information Technology Standards, DOI 10.1007/978-3-319-30370-3_7

replete with homonyms (where the same term means different things depending on context), synonyms (where there is more than one term for exactly the same concept), eponyms named after people, three letter acronyms and abbreviations. Nobody, who has not learnt the eponym, can guess the meaning of `Hodgkin's` (lymph node cancer), `Bright's` (kidney disease) and `von Recklinghausen's disease` (hereditary neurofibromatosis).

Information scientists classify knowledge in a series of levels. For example the Dewey Decimal Classification, used in libraries, attempts to organize all knowledge into ten main classes:

000 – Computer science, information and general works
100 – Philosophy and psychology
200 – Religion
300 – Social sciences
400 – Language
500 – Science (including mathematics)
600 – Technology and applied Science
700 – Arts and recreation
800 – Literature
900 – History, geography, and biography

Blois in his seminal book *Information and Medicine* showed how healthcare is unique amongst scientific endeavors in that day-to-day medical information relates to so many different levels [1]. The breadth of healthcare covers an exceptionally broad scope, ranging from radiation and subatomic structures, complex molecules including DNA and proteins, cells including hematology and cancers, micro-organisms such as bacteria and viruses, anatomical structures including the different body systems, mental activity, the whole person, groups, societies and populations.

Each aspect of healthcare mixes multiple overlapping theories, each with their own sub-terminology. Any classification system is inevitably just one way of slicing up a very complex reality, made even more difficult because key medical concepts such as diseases are abstractions, defined using information from a variety of information levels; diseases are not objects which can be seen or touched.

People use terms in the way that they and their immediate colleagues understand. Each user of a term assumes that everyone else understands precisely what he or she intends it to mean; over time groups develop their own local dialect. Medical records staff can often identify not only a doctor's specialty but also the institution where he or she was trained from the way they use certain terms.

Lewis Carrol expressed the same problem in an exchange between Alice and Humpty Dumpty in *Through the Looking Glass*: [2]

> *'I don't know what you mean by "glory"' Alice said.*
>
> *Humpty Dumpty smiled contemptuously. 'Of course you don't – till I tell you. I meant "there's a nice knock-down argument for you!"'*
>
> *'But "glory" doesn't mean "a nice knock-down argument"' Alice objected.*
>
> *'When I use a word,' Humpty Dumpty said in a rather scornful tone, 'it means just what I chose it to mean – neither more nor less.'*

'The question is,' said Alice, 'whether you can make words mean so many different things.'

'The question is,' said Humpty Dumpty, 'which is to be master – that's all.'

The representation of written information has become more and more specific over the centuries. The first way of representing information was by a picture or drawing, such as in Stone Age cave paintings. The earliest writing was based on pictograms, such as Egyptian hieroglyphics and Chinese characters, but the need for cheap and quick writing materials led to the development of cruciform characters on wet clay blocks in Mesopotamia and the development of phonetic alphabets such as those of Greece and Rome. Modern English does everything using just 26 letters, 10 digits and a few punctuation marks.

Computers hold information as sequences of binary bits (`0` s and `1` s) and work by matching strings; they need precisely coded data. A computer can instantly check if two strings are the same but, if a difference is detected, it cannot judge whether that difference is important. In spite of decades of effort we do not yet have computers that cope well with the ambiguity inherent in natural language.

Coding and Classification

People often confuse the terms coding and classification and use them almost synonymously. This may be because the process of classification involves recording the codes used to name specific classes. However, coding schemes and classifications do different jobs. Classification allocates things into groups or classes, while coding is the allocation of identifiers, which can apply to anything (including classes in classifications).

A **code** is a sequence of symbols, usually digits or letters, which designate an object or concept for identification or selection purposes. It is simply an alternative name for something, an identifier, designed for computer processing. Coding systems are an indispensible part of healthcare computer applications and interoperability specifications for exchanging data between computers.

The primary challenge for the designers of coding schemes is to produce something that will be widely and willingly adopted and endorsed by clinicians and managers. However, clinicians and managers have no more interest in codes than a retail customer has in the bar code on a packet of corn flakes. Codes are needed and used by computers, not humans.

Clinicians need to record information in the form, language and detail that is of most benefit to them when treating individual patients. Clinical records require precise and comprehensive detail about each individual patient, creating a tension with statistical analysis, which requires patients to be classified into a manageable number of discrete and mutually exclusive groups.

Clinicians and managers should to be interested in classification, because it is the basis for most statistical analysis, quantitative management, accountancy and research.

Classification is the systematic placement of things or concepts into categories or classes, which share some common attribute, quality or property. There is no limit to the number of ways that any set of objects can be classified and so no possibility of a perfect classification that is good for everything.

The choice of what classification system to use is often determined by payment agencies, insurance companies and national governments that control whether or not a doctor or institution gets paid. Such bodies usually specify the precise classification system that they require, often in collaboration with representatives from the professional and trade associations, medical colleges and educational bodies. Once chosen it has to be accepted by users and implemented in computer software.

In *The Endangered Medical Record*, Slee argues that the choice of scheme used for electronic patient records represents a serious real threat to the truthfulness and completeness of medical record content. By using of broad categories, such as those specified by the International Classification of Diseases (ICD), rather than precise diagnoses, we throw away detail that should be preserved permanently. His plea is for detailed, permanent and unambiguous codes [3].

For example, a trauma surgeon might describe a typical skiing accident as: a `closed spiral fracture of the shaft of the right tibia with fractured fibula`. In ICD-10, the code for fracture of shaft of tibia has the following logical structure:

```
Chapter XIX: Injury, poisoning and certain other consequences of exter-
nal cause (S00-T98)
    Block: Injuries to the knee and lower leg (S80-S98)
      S82: Fracture of lower leg, including ankle
             S82.2 Fracture of shaft of tibia (with or without mention
             of fracture of fibula)
                          S82.2.1 Closed fracture of shaft of tibia
```

The selected ICD-10 code `S82.2.1` does not specify whether the leg is left or right, whether the tibia fracture is simple, spiral or compound or how the fibula is affected.

A **hierarchy** is an ordered organization of concepts. General concepts are at the top of the hierarchy; at each level down the hierarchy, concepts become increasingly specialized. This can be thought of as an inverted tree with its trunk or root at the top. For example, biological classification places animals and plants into a hierarchical classification (a taxonomy) according to similarities in structure, origin etc., which indicate a common relationship. The main levels in the biological taxonomy are Kingdom, Phylum (animals) or Division (plants), Class, Order, Family, Genus and Species.

Healthcare computer systems use nationally prescribed coding systems. Many of these, such as the ICD-10, CPT-4 and Read Codes use a *position-dependent hierarchical coding structure*. The internal structure of the code specifies its meaning relative to other codes. The structure of the code increases in detail from left to right, with the first character of the code specifying the chapter, the second the main subdivision and so on until down the branches of the tree until the final leaf codes are reached.

One of the technical problems of position-dependent hierarchical coding systems is that they cannot be modified easily without changing the meaning of codes in different versions, creating major problems when, as inevitably happens, one version needs to be replaced by another.

Coding Systems

Any coding system has various components.

Concept: The fundamental idea is that of a concept, which is a medical idea. Each concept is identified by a concept code.

Coding Scheme: Each concept code originates from a coding scheme. A coding scheme defines a set of concept codes, which are unique within the namespace of the coding scheme, and are globally unique when coupled with the name of the coding scheme itself.

Display Term: This is a human readable term. In some cases more than one display term may be provided for the same concept, to cover true synonyms, such as translations into different languages. One display term is usually designated as the preferred term.

Relationship: Concepts may be related to other concept via a relationship, which allows the generation of hierarchical structures. One concept may be part of more than one hierarchical structure. Often these relationships will be defined as part of original coding schemes, but other relationships are also possible.

Value Set: A set of values that are allowed for a particular data item. Message specifications refer to value sets as the allowed values for a field. Codes from a single coding scheme may be referenced using a value set table, which has a heading and includes metadata such as: value set name, unique identifier, coding scheme, author, time validity, version and other notes. Each entry in the table contains concept code value, display term and notes about applicability.

Identifiers: Computer systems need unique identifiers for people, things and places, which have similar properties to codes. One way of achieving uniqueness is to treat each identifier as a pair, comprising a unique name for the assigner plus a value for the identification number, which is unique within assigner. It is the responsibility of the assign or to ensure that all such values are unique.

Terminologies

Terminology: a set of concepts designated by terms belonging to a special domain of knowledge, or subject field.

Reference terminology: a terminology in which every concept designation has a formal, machine-usable definition supporting data aggregation and retrieval. Reference terminologies are designed to provide exact and complete representations of a given domain's knowledge, including its entities and ideas, and their interrelationships, and are typically optimized to support the storage, retrieval, and classification of clinical data.

Interface Terminology: Systematic collections of clinically oriented phrases or terms aggregated to support clinicians' entry of patient information directly into computer programs, such as clinical documentation systems or decision support tools. They may mediate between a user's colloquial conceptualizations of concept descriptions and an underlying reference terminology.

Ontology: hierarchical structuring of knowledge about things by sub-categorizing them as a set of concepts within a domain according to their essential qualities and relationships between those concepts.

Expression: A collection of references to one or more concepts used to express an instance of a clinical idea. An expression containing a single concept identifier is referred to as a pre-coordinated expression. An expression that contains two or more concept identifiers is a post-coordinated expression.

The scope and some of the terms used in clinical terminology are summarized in Fig. 7.1.

User Requirements of Terminologies

A key design requirement for any coding and classification system is to satisfy the needs of the different stakeholders. Roger Côté, the father of SNOMED, views this as a pyramid with three levels of use:

1. At the tip, case-mix classifications such as DRGs, used for payment.
2. In the middle, classifications of diagnoses and procedures used to monitor and audit clinical activities.
3. At the base, clinical terminology used for individual patient care.

Healthcare managers and researchers need classified data, which enable comparisons and data exchange with existing data sources. Links between classifications must be explicit with one-to-one or many-to-one links. A many-to-one link involves loss of information, the extent of which is determined by how closely one classification is based on the other.

A multilevel classification with both coarse and fine granularity may allow two-way mapping from another classification. High levels of compatibility can usually

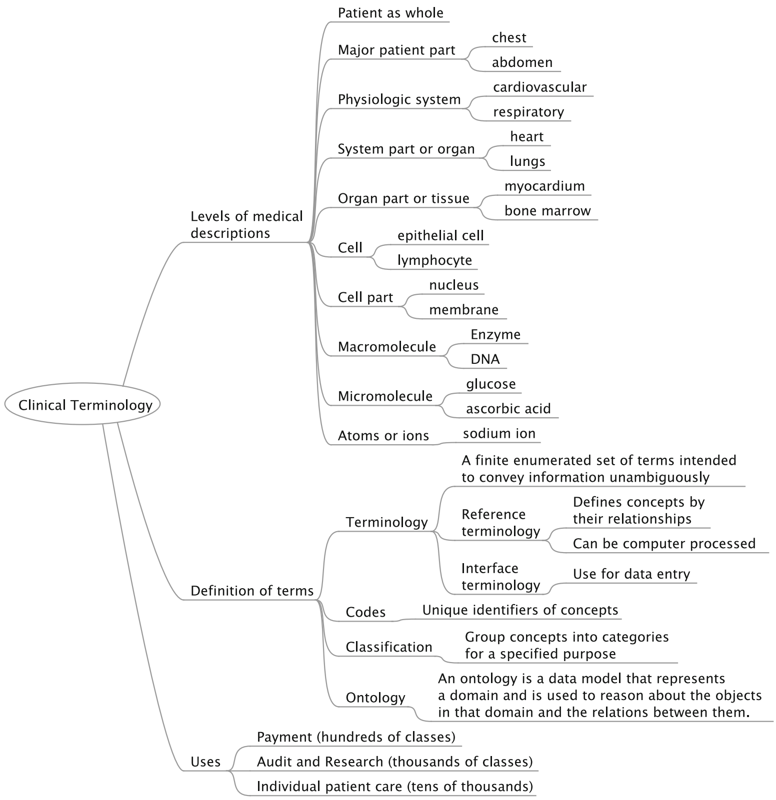

Fig. 7.1 Scope and terms used in clinical terminology

be obtained only by basing a new classification directly on the target, using the same class boundaries. This requirement for cross-mapping with existing classifications inevitably drives the developers of clinical classifications to build on existing schemes, even if they are not suitable for the need in hand. For example the ICD is organised around body systems, which is helpful in some circumstances, but not in others. Early versions of SNOMED reflected its origins in the College of American Pathologists as an extension of the Systematized Nomenclature of Pathology (SNOP), which gave a pathological slant.

Doctors and nurses will not take the trouble to learn how to use any system unless it is quick and easy to use and provides information in the form and language that best helps them treat individual patients. Automatic or semi-automatic encoding software is needed. Clinical records need to be as specific as possible. Hence clinicians require a comprehensive nomenclature of medical terms covering everything that could occur within any patient's medical record. That is, all of clinical medicine and health service administration, but not the whole of bio-medical science.

In 1984, the IMIA working conference on clinical terminology concluded that

In future healthcare information systems, the user interface should be based upon natural language. The generation of numerical or alphanumeric codes should occur within the computer. Automatic encoding of natural language should be used. The morbidity and mortality statistical classification requirements of national and international groups should be the by-product of medically based healthcare information systems.

It was not anticipated that almost 30 years later most clinical coding in hospitals would continue to be done by coding clerks.

Desiderata

Desiderata for Controlled Medical Vocabularies in the Twenty-First Century [4] brought together together a number of common requirements for clinical terminologies, which had been developed in leading terminology projects such as GALEN [5], UMLS (Unified Medical Language System) [6], SNOMED RT (Reference Terminology) [7] and the NHS Clinical Terms Project [8]. This paper was highly influential in the design of SNOMED CT. The desiderata are:

1. **Vocabulary Content:** In terms of scope and quality, content is paramount. Any practical clinical terminology needs to be comprehensive in terms of both domain coverage (concepts) and human readable terms (descriptions and synonyms). A methodology is required to allow the content to be expanded as and when required, including translation into other languages and dialects, while maintaining quality.
2. **Concept Orientation:** This means that each concept term has one meaning (non-vagueness), and only one meaning (non-ambiguity). However, each concept can be described by several terms (synonyms) in the same language plus different terms in other languages and dialects. Note also that the same term can have different meanings (homonyms), each relating to a different concept.
3. **Concept Permanence** Once a concept is created its meaning persists. It must not be changed or deleted by updates. However, a concept may be marked as *retired* where its meaning is found to be ambiguous, redundant or otherwise incorrect.
4. **Non-semantic Concept Identifiers** Each concept should have a unique identifier, which should be meaningless. All semantic information (relating to meaning) is an attribute of the concept, and should not be part of its identifier. Further examples of the problems with position-dependent hierarchical coding schemes are discussed with examples in Chap. 8.
5. **Polyhierarchy** While it is useful to organize medical concepts in a hierarchical way, many clinical concepts are naturally multi-dimensional, with more than one parent (super-type) concept. For example, a *fractured tibia* is both a type of *fracture* and a type of *leg injury*.
6. **Formal Definitions** The means of classifying a concept is independent of the means of identifying it. The development of formal, descriptive logic to define and classify clinical concepts is a major development away from the traditional position-dependent coding schemes and dictionary forms of definition. For

example *pneumococcal pneumonia* may be defined using a hierarchical (*is a*) link to the concept *pneumonia* and a *caused by* link to the concept *streptococcus pneumoniae*.

7. **Rejection of "Not Elsewhere Classified" Terms** Many existing classifications include one or more catchall categories for concepts not covered. The problem with such *not elsewhere classified* or *NEC* categories is that they change their meaning, as and when a new category is added that covers some of the NEC scope. The meaning is not permanent, which was a previous criterion.
8. **Multiple Granularities** Different users require different levels of granularity. Different levels of granularity are needed for defining concepts, navigation, decision support and reporting. For example, a manager may only need to know that a patient has a *broken leg*; the finance department that it is a *fractured tibia*, but the clinician needs to know that it is a *closed spiral fracture of the shaft of the right tibia*. In principle, there should be no limitations on the number of levels in the display tree hierarchy.
9. **Multiple Consistent Views** When a concept has multiple parents in a hierarchy, the view of that concept should not depend on whether it was reached by following the hierarchy from a particular parent. The complete structure of a terminology, including all hierarchies and relationships can be complex and difficult to use. Each end user needs one or more views that reflects his or her own needs and understanding, but in a way that is consistent with the underlying model.
10. **Context Representation** Information is recorded within a particular context and cannot be interpreted without that understanding. The context needs to be computer-processable. One approach is to provide a means of recording context explicitly within the terminology.
11. **Graceful Evolution** Terminologies change over time. It creates problems for users if the meanings of aggregated time series data change in an uncontrolled manner. Care is needed to design the whole structure to support graceful evolution of concepts, terms and relationships.
12. **Recognize Redundancy** When terminologies change, some components will become redundant and so it important to recognise explicitly that this has happened.

The Chocolate Teapot

The apocryphal story of the Chocolate Teapot, developed by Dr Malcolm Duncan [9] illustrates a number of the issues involved in classification and terminology. I am grateful for his permission to reproduce this in an edited form.

Because not all readers will be confident they know what, for example, asthma really is (even doctors disagree), this discussion is organised around the classification of teapots. Most people think they know what a teapot is, or do until they read this.

Consider a fragment of a simple but functional crockery classification that exists solely to document the appearance of table settings as described in Victorian literature.

Classification of Tableware Version 9 (CTV-9, 1875)

```
Crockery
---Teapot
------Brown teapot
------White teapot
------Blue teapot
------Teapot, color not elsewhere classified (NEC)
------Teapot, color not otherwise specified (NOS)
```

Users are obliged to *code to the leaf*, which is how ICD-9 and ICD-10 are used. Only brown, white and blue crockery were fashionable in 1875. The uncouth might deploy other colors (`NEC`) or worse, might not care (`NOS`). However it is axiomatic that your teapot is made from earthenware, ideally from quality porcelain.

Elsewhere there is the frequently updated *Systematized Nomenclature of Kitchen Terminology (SNoKitch)* intended to support all catering applications. A fragment of this terminology is followed through a number of iterations:

SNoKitch Release n

```
Crockery
---Teapot
```

Teapot has no children and could be equivalent to any of the five classification leaves of CTV-9. People then ask for further concepts to be added, leading to the next release.

SNoKitch Release n + 1

```
Crockery
---Teapot
----- Brown teapot
------White teapot
```

Anyone wanting to specify a brown or white teapot using *SNoKitch* can now do so. Coding to leaf is not mandatory and so `NOS` and `NEC` are absent. The Teapot concept in *SNoKitch* is the equivalent of both `Teapot, color (NEC)` and `Teapot, color (NOS)` categories in CTV-9, although in effect the meaning of `Teapot` in this version of SNoKitch is skewed to mean a teapot that is neither brown nor white.

SNoKitch Release n + 2

```
Crockery
---Teapot
```

```
----- Brown teapot
------White teapot
------Blue teapot
------China teapot
```

In release n+2 `Blue teapot` (forgotten in the previous version) is added as a child. This further alters the meaning of the concept `Teapot` in its role as `Teapot, color not elsewhere classified`.

Also a `China teapot` child has been added, but this is not a discrete sibling: it could be any color. `Teapot` now has non-disjoint subclasses. `China teapot` has perhaps acquired part of the meaning of `Teapot, color (NEC)` from `Teapot` (ie China teapots which are not brown, white or blue). However, when users interested in whether a teapot is china receive a `White teapot` code they are out of luck. All they can infer is that it represents a teapot.

As both the material and color are important, new concepts `White china teapot` and `Blue china teapot` are added in the next release. Unfortunately we cannot rely on everyone using them because they don't have to code to leaf. Some people may only choose to specify whether their teapots are china and not capture color at all.

However we still have a common understanding of the meaning of the parent `Teapot` concept as superclass of its children. Concepts added in the next release will overturn this.

SNoKitch Release n + 3

```
Crockery
---Teapot
----- Brown teapot
------White teapot
----------White china teapot
------Blue teapot
----------Blue china teapot
------China teapot
---------White china teapot
---------Blue china teapot
------Chocolate teapot
------Ornamental teapot
------Industrial teapot
```

The addition of `chocolate`, `ornamental` and `industrial` teapots means that ontological continuity with previous versions is deeply in question. `Industrial teapot` indicates where it is used but nothing else except presumably that it is not ornamental or made of chocolate. Teapots made from metal or chocolate are not earthenware and hence not crockery. From release to release, there is little consistency in what we can infer about color, material or use.

Worse, the arrival of certain children has even altered the meaning we can infer from the unadorned `Teapot` parent concept. For example, the *Kansas Tea Company (KTC)* operate decision support software designed against release n + 2:

```
If Concept = (teapot or child of teapot)
Safe to add tea leaves plus boiling water
End If
```

KTC had not anticipated confectionary in this hierarchy. As for `ornamental teapot`, who knows if it can be used to make tea? KTC must now add multiple additional nodes to their decision support rule and be able to exclude subtypes eg it would not be wise to add boiling water to `Chocolate teapot`.

The validity of these relationships depends on your interest in teapots. There is no longer a universal understanding of the concept. There is merely a shared assumption about the *term* now exposed as not applying in all contexts of use. If we had started with a clear definition of teapot (as a 'free text' scope note) this might have been avoided. Most people can express the teapot that is in their head most of the time, but this may be at the expense of loss of predictable machine readability.

The *Comestibles Supply Consortium* recognises the incoherent use of the *SNoKitch* terminology across their systems and mandates use of a small subset.

```
Teapot
------Brown teapot
------White teapot
------Blue teapot
```

All within the Consortium will now understand what was recorded and satisfy their leading concern which is to make billable returns using *Classification of Tableware Version 9*. They still lack the distinction between NOS and NEC and may need to employ professional coders to abstract and map records manually.

Classifications such as the ICD family include items like `Chronic airway obstruction, not elsewhere classified` and `Other specified excision of adrenal gland`. A problem with such constructs is that they are not stable in meaning across versions of the classification ie what is classified elsewhere may change with addition or removal of other content. In contrast to classifications such as ICD-9, modern medical terminologies such as SNOMED CT typically do not mandate coding to leaf and do not permit *not otherwise specified* or self-referential entities such as *not elsewhere classified*.

The widespread assumption is that this provides immunity from semantic discontinuity across releases. It does not. In the absence of NOS and NEC, these static 'known unknowns' become 'unknown unknowns' mobile between releases. To quote Donald Rumsfeld:

> *…there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns — the ones we don't know we don't know.*

The situation is further exacerbated if a terminology is neither exhaustive at any one *level* of siblings nor disjoint (two sets are said to be disjoint if they have no element in common). As terminologies evolve, applications involving data reuse (such as messaging, billing, clinical audit and active decision support) need to recognise what version of the terminology a code is drawn from. Successive versions may improve content but these alterations create circumstances where interpretation of recorded data requires reference to the ontology, as it was when the concept was chosen, not as it is now.

A concept within a terminology is not entire unto itself. Addition, retirement and movement of other concepts alter its use and interpretation. Such changes are common. As well as additions and retirements there may be many hierarchy changes in each release of a terminology such as SNOMED CT.

References

1. Blois MS. Information and medicine: the nature of medical descriptions. Berkeley: University of California Press; 1984.
2. Carrol L. Through the looking glass and what alice found there. London: Macmillan; 1871.
3. Slee V, Slee D, Schmidt HJ. The endangered medical record: ensuring its integrity in the age of informatics. St. Paul: Tringa; 2000.
4. Cimino JJ. Desiderata for controlled medical vocabularies in the twenty-first century. Methods Inf Med. 1998;37:394–403.
5. Rector A, Solomon W, Nowlan A, Rush T, Claassen A, Zanstra P. A terminology server for medical language and medical information systems. Methods Inf Med. 1994;34:147–57.
6. Lindberg DA, Humphreys BL, McCray AT. The unified medical language system. Methods Inf Med. 1993;32(4):281–91.
7. Spackman K, Campbell K, Cote R. SNOMED RT: a reference terminology for healthcare. Proc AMIA Symp. 1997: 640–4.
8. O'Neil MJ, Payne C, Read JD. Read codes version 3: a user led terminology. Methods Inf Med. 1995;34:187–92.
9. Duncan M. Medical terminology version control discussion paper: the chocolate teapot. Version 2.3. Medical Object Oriented Software. 2006. www.mrtablet.demon.co.uk