Introduction of a Clinical Terminology in The Netherlands

Needs, Constraints, Opportunities

NICTIZ

Nationaal ICT Instituut in de Zorg

postadres: Postbus 262, 2260 AG Leidschendam
bezoekadres: Overgoo 11, 2266 JZ Leidschendam
telefoon: (070) 317 34 50; fax: (070) 320 74 37; e-mail: info@NICTIZ.nl

Versie : 1.0
Datum : September 2003

Authors:
Pieter E. Zanstra, Egbert J. van der Haring, Ronald Cornet[*]
Depts. Medical Informatics, UMC St Radboud Nijmegen, [*]AMC Amsterdam
Contact: p.zanstra@mi.umcn.nl

Contents

Samenvatting

De zorg is een zeer informatie-intensieve sector. Dit geldt voor zowel de directe patiëntenzorg alsook voor de activiteiten voor registratie en declaratie. Het gaat in de zorg niet alleen om financiële verslaglegging, waar we de euro als registratie-eenheid hebben, maar vooral ook over de vastlegging van allerlei gegevens over de patiënt. Voor dat laatste bestaat geen eenduidige universele taal. In de huidige situatie wordt gebruik gemaakt van classificatiestelsels welke bedoeld zijn om gegevens van patiënten in categorieën van bijvoorbeeld ziekten of verrichtingen in te delen. Deze stelsels zijn vaak niet ontworpen voor directe patiëntenzorg maar voor statistiek en verrekening van kosten, en daarmee a-priori ook minder geschikt voor de vastlegging van het dagelijks handelen. De stelsels missen het detail dat nodig is in de directe zorg, en voor een veilige gegevensuitwisseling. De door velen gepropageerde eenmalige vastlegging in één enkel (virtueel) patiëntendossier, welke voor alle geautoriseerde behandelaren toegankelijk is, vereist dat termen niet alleen op eenduidige wijze worden vastgelegd maar ook door anderen (medebehandelaren) op dezelfde wijze worden begrepen als ze door de vastlegger bedoeld zijn. Dit vraagt om eenheid van taal.

In het afgelopen decennium is er een groeiende belangstelling voor 'klinische terminologie' waar te nemen. Het begrip kan omschreven worden als: *De verzameling van standaard termen met hun synoniemen, die in de directe patiëntenzorg gebruikt kan worden voor de vastlegging van alle klachten, symptomen,omstandigheden, ziekteprocessen, interventies, diagnosen, resultaten, en de besluitvorming*. In een samenwerking tussen de Engelse NHS (National Health Service), en CAP (College of American Pathologists) is in de afgelopen jaren een klinische terminologie ontwikkeld en op de markt gebracht onder de naam SNOMED CT (Clinical Terms).

Hoofddoel van het onderhavige onderzoek is vast te stellen of het zinvol is om op korte termijn tot een brede invoering van SNOMED in Nederland te komen, en zo ja voor welke specifieke zorgsectoren. Secondair is de vraag of er alternatieven zijn, en wat de kosten van dergelijke opties zijn. Tijdens dit onderzoek hebben we dankbaar gebruik kunnen maken van de vele publieke, maar ook interne documenten die ons met name vanuit het Verenigd Koninkrijk en Australië ter inzage zijn verstrekt.

Een zorgbrede invoering van een product als SNOMED betekent, zeker wanneer men de meerwaarde wil gebruiken, een revolutionaire verandering in het proces van verslaglegging. De vraag is of dat nodig, wenselijk, of zelfs haalbaar is. Een evolutionaire invoering lijkt vooralsnog meer de aangewezen weg.

Geconcludeerd kan worden dat, hoewel SNOMED de voornaamste kandidaat is om de de facto wereldstandaard te worden, het nog niet voldoende rijp is. Tegelijkertijd moet ook vastgesteld worden dat er nog een fundamenteel gebrek aan inzicht is in de behoeften ten aanzien van een geavanceerde terminologie in de zorgsector. Dit vloeit mede voort uit onduidelijkheden over de overlap tussen terminologie en dossierarchitectuur. Aanvullende experimenten, waarbij SNOMED een belangrijke rol kan spelen, worden sterk aangeraden, om eisen aan terminologie helder op tafel te krijgen. Aangezien het verhelderen van deze eisen in samenhang met het gebruik in dossiersystemen de belangrijkste prioriteit is, is onderzoek naar de alternatieve terminologieën op dit moment minder relevant. In een appendix is enige publieke informatie over die alternatieven opgenomen.

Hoewel internationaal, vooral in ICT kringen, de recente discussie over klinische terminologie suggereert dat deze alleen in dossiersystemen gebruikt wordt, moet hier toch met klem op secondair gebruik van die gegevens gewezen worden. Belangrijk gebruik is er ook voor statistiek, epidemiologie, indexeren en zoeken van literatuur, en verrekening van zorg. Het aantal factoren in het veld, de toenemende complexiteit, en onderlinge afhankelijkheden, vereisen een meer geïntegreerde benadering van terminologie gerelateerde zaken.

In hoofdstuk 8 van dit document worden een zestal aanbevelingen gedaan. In het kort worden ze hier samengevat:

1. Voer tenminste een tweetal substantiële 'real life' experimenten met SNOMED uit. Besluit op basis daarvan over eventuele invoering van SNOMED als nationale standaard.
2. Ga door met een meer op formele principes gebaseerd onderhoud en ontwikkeling van de bestaande classificaties in de zorg. Onderzoek ook of deze bestaande classificaties als een startpunt voor ontwikkeling van een eigen klinische terminologie kunnen dienen.
3. Richt een nationaal expertise centrum voor zorg gerelateerde terminologieën op.
4. Zorg voor continu verzekerde toegang tot gecodeerde zorggegevens. Houd zoveel mogelijk vast aan open source condities voor een nationale klinische terminologie.
5. Ontwikkel criteria voor de specificatie en certificatie van terminologieën en afgeleide producten
6. Wees er op voorbereid dat er voor zorgterminologie jaarlijks een centraal budget van tenminste 2 miljoen Euro nodig zal zijn.

Executive Summary

Health language is still far from being harmonised. Present day classifications like ICD and procedure classifications have not been designed for use in direct individual patient care. The purpose of the registration of data is to have a faithful record of what is done with or for the patient. Simultaneous or consecutive medical attention to the patient requires insight and understanding of all participants in the treatment. This is a very ambitious goal, but the change of healthcare delivery from the individual carer to care teams dictates this change in recording habits. The often-postulated single virtual patient record that is accessible for all authorised carers is implicitly based on this ambition. The central requirement is that everyone uses and equally well understands the common language used in the record.

The main objective of this study was to answer the question 'Is it sensible to aim for a national introduction of SNOMED CT, and if so for which sectors of healthcare?' The secondary questions were on alternatives and an estimate of costs involved.

In collaboration with the National Health Service (NHS), the College of American Pathologists (CAP) has developed a very detailed clinical terminology (SNOMED CT). The full introduction of a product like SNOMED means a revolutionary change in recording practice. The question is whether this is necessary, desirable, or even feasible. An evolutionary process of introducing technologies is certainly worth considering.

This study is based on interviews and numerous internal documents that were kindly shared by the interviewed persons, as well as documents from public resources.

The conclusion on SNOMED is that, though it is the leading candidate to become a defacto global standard, it has not yet reached sufficient maturity. At the same time it is concluded there also is a fundamental lack of understanding on terminology requirements in the (Dutch) healthcare sector. Further research is strongly recommended. In experiments SNOMED can serve as an instrument to elicit ate those requirements. As long as the user requirements are not sufficiently clarified a further investigation of alternative terminologies is considered as not yet relevant. A few prototypical Open Source and commercial examples are described in some detail in the appendixes.

Though recent developments suggest that terminology is only used in the context of the health record, it should be observed that there are many secondary usages of coded data such as in health statistics. The number of actors in the field, the growing complexity and interdependencies, do demand a more integrated approach of all matters of health terminology.

The recommendations (Chapter 8) should be read in their context. In short they are:

1. Do conduct at least two major real life learning experiments with SNOMED. Make on the basis of the outcome a decision on SNOMED as a national standard
2. Continue with a more principled further development and maintenance of existing reporting classifications. Do investigate if this new approach is a viable start for a clinical terminology.

3. Establish a national expertise centre for all Health Terminology
4. Secure appropriate access to coded data. Insist on sufficiently Open Source conditions for a national reference terminology
5. Further develop criteria for specification and certification of terminologies and derived products
6. Be prepared for a central annual expenditure of at least 2 Million Euro on health terminology

1 Motivation

The healthcare delivery process is an extremely information intensive industry. This is the case for direct patient care as well as for registration and billing activities. Because the healthcare process takes a substantial part of the national budget, the society demands to know a lot of that process. Where the financial industry has the currencies as a relatively simple common ground, health language, the 'currency' of the health care delivery process, is much more complex, and still far away from being harmonised. To achieve unambiguous recording of patient data across the sector, a good common understanding of the words used is very important. Many have expressed the need for terminology that can be understood by all either directly, or after being 'translated'.

Direct patient care requires rather detailed terminologies to closely describe the patient's condition. Most of present day classifications like ICD and procedure classifications have not been designed for use in direct individual patient care. They were primarily designed to describe at a higher level and for purposes of counting. Clinicians do not consider the language of existing classifications as their own language. This results in the development of many local lists and adaptations of varying quality and compatibility.

Ideally for any patient case there ought to be a detailed recording of findings, diagnostic procedures with results, diagnoses, and the reasoning that has led to such conclusions. Also treatment and results, as well as any acts that have been performed as part of the treatment plan need recording. This must be done in an unambiguous language within a structure that makes it reusable and useful for any qualified healthcare professional at any time, at any place. In that sense there is in essence no difference between sending patient data from one ward to another or for transfer between shifts. Current practice shows that relatively little is recorded in a standardised way. Secondary care has just barely started with the introduction of a real electronic health record. Quite often one sees free text from which it is hard to reconstruct clinical detail. Some documents are even stored as scanned images, which render them totally useless for automatic processing. Laboratory data and drug prescriptions generally are reported in a more formal way. The same is the case for discharge diagnoses for national registries, and procedures for reimbursement purposes. In some specialties there are national or international groups that have agreed on a common registry for research purposes.

The purpose of the registration of data is to have a faithful record of what is done with or for the patient. This recording must be useable for multi-disciplinary treatment of the patient. Also the increase of part-time doctors dictates that more doctors of the same specialty must be able to communicate about the same patient via one single (virtual) record. Simultaneous or consecutive medical attention to the patient requires insight and understanding from all participants in the treatment. This is a very ambitious goal, but the change of healthcare delivery from the individual carer to care teams dictates this change in recording habits. The often-postulated single virtual patient record that is accessible to all authorised carers is implicitly based on this ambition! The central requirement again is everyone uses and equally well understands the same language.

In the past decade we have seen a growing interest in what is called 'clinical terminology'. This term stands for the terminology needed to record all the important elements of patient care in standardised terms. This assumes that for a patient encounter significantly more will be recorded than done in current practice. In principle

such a terminology is supposed to represent 'all of medicine'. If that ultimately will be one single resource, or a collection of compatible terminologies is a matter of growing debate. There is not an internationally accepted definition for 'clinical terminology'. It has been described as:

A clinical terminology is the collection of standard terms with their synonyms, which in the context of patient care support the recording of complaints, signs, symptoms, circumstances, process of illness, interventions, results, diagnoses, as well as the decision making of the care providers.

Descriptions in these standard terms should map unequivocally to categories in the existing (inter-) national classifications for diseases, procedures, reimbursement etc. Some argue that there should ultimately be one single global reference to support all kinds of multidisciplinary care pathways.

In collaboration with the National Health Service (NHS), the College of American Pathologists (CAP) has developed a very detailed clinical terminology (SNOMED CT) which at the time of this writing contains about 350.000 concepts, which are linked to some 1.5 million textual English and Spanish expressions. In both the UK and the USA there are national plans for using this language for recording. Internationally SNOMED CT is seen as the emerging de facto standard for clinical terminology. Within HL7 it is widely considered as the most important standard terminology for communication. Compliance with these emerging standards offers both opportunities and threats to the Dutch healthcare IT industry. SNOMED is not incontestable. There are questions about a too high level of detail, high costs of deployment, influence of the local user on the product, and ownership of local additions of the terminology, to name just a few concerns.

The central focus of NICTIZ is to develop an infrastructure that enables access by authorised care providers to all relevant patient data independent of time and place. In principle this means that every healthcare information system should 'add' information to the virtual electronic record. Can or should SNOMED CT be deployed as the common language to facilitate interdisciplinary information exchange?

The full introduction of a product like SNOMED means a revolutionary change in recording practice. The question is whether this is necessary, desirable, or even feasible. Daily practice shows time and again that changes occur much slower than anticipated. In different healthcare sectors developments occur at different speeds. Consequently an evolutionary process of introducing new technologies is likely to be the only option.

2 Objectives

The objective of this study is to make a number of recommendations on deployment of clinical terminologies in The Netherlands for the period 2003-2008. Originally, the formal task of this study was to answer the question: "What is the added value of SNOMED CT with respect to improvement of communication and interoperability between information systems in healthcare". Early in the project it was decided that this assignment was too narrow. The focus should not just be on information systems, but on the healthcare process as a whole.

It is difficult to define a priori what the size and level of detail of a clinical terminology should be. To what extent can one cope both from a cognitive as well as a technical point of view? Who are the stakeholders? Do we need one clinical terminology for all of medicine, or is there one central core with extensions per specialty. How do we manage the ever-changing health language? How do we deal with conflicting definitions? There are many questions to be answered. What matters now is to determine the most immediate questions, and those that can be postponed.

For this study we deliberately choose not to conduct an intensive needs assessment for a rich detailed clinical terminology. First of all, we did not have the time for it, and secondly such a study would be costly because of the size and the diversity of the healthcare domain. There is also no immediate need for such a study because there are large resources available that can be used in evaluation studies. One of the most obvious choices is SNOMED CT, which is the most visible terminology in this field. Given the existence of such a large structured terminology we can design experiments to determine the conditions for deployment of a large terminology in healthcare at a fraction of the development cost.

The project plan seeks to answer the following questions:

1. What are the major problems with respect to usage of terminology in the primary process of healthcare as a whole (primary, secondary, and tertiary care and cure)?

2. SNOMED CT and some other (commercial) products position themselves as the solution.
   a. Which are those other products
   b. What are the essential characteristics of those other products

3. As stated above, the emphasis of this study is on SNOMED CT. Is SNOMED potentially the solution, and if so:
   a. For which application areas
   b. (How) can SNOMED be implemented in the existing Hospital Information Systems (legacy)
   c. What is the expected impact for the care provider for using this new terminology
   d. Does it provide sufficient detail for documenting the care process

e. What does the usage of SNOMED mean with respect to maintenance. (How is the editorial control of the Dutch version, who is responsible for local adaptations, who is the owner of the adaptations, etc.)
f. Will the introduction follow an evolutionary pathway, or is a 'big-bang' revolutionary introduction necessary
g. What is the estimation of costs with regard to:
    i. Licensing
    ii. Porting to the Dutch situation
    iii. Maintenance and Training

This study is based on interviews and numerous internal documents that were kindly shared by the interviewed persons, as well as documents from public resources.

The outcome of the study is an inventory and decision-making document on the possible introduction of SNOMED CT as the common language for healthcare. On the basis of interviews and the background documentation preliminary statements are expected on the following:

- Recommendation to NICTIZ/ The Dutch Healthcare sector

Is it sensible to aim for a national introduction of SNOMED CT, and if so for which sectors of health care?

or

Is it sensible to further study other (commercial) alternatives?

Is it sensible to start a new development for a Dutch clinical terminology?

- Estimates of costs involved with the introduction of a clinical terminology for the whole healthcare sector on the basis of:

Adoption of SNOMED CT

Development of a Dutch terminology

Other (commercial) alternatives

- Recommendations with respect to usage and/or extensions of existing classifications and coding systems

- A strategy for a possible (stepwise) implementation of a clinical terminology (migration)

3 Problems with health terminology

Currently we are entering the third phase of introducing information technology to the health care delivery system. The first phase was mainly concerned with automating the administrative process. The second phase was about delivering medical applications. The third phase is about integrating a diversity of medical and administrative systems in one coherent interoperable environment.

At present we see an emerging need for safe and sensible communication between these applications. Next to a good insight of what is really necessary to exchange, also the language used is important. For instance, the slow progress in the application of knowledge-based systems may to a great deal be attributed to the lack of semantic coupling with the patient record.

Medicine is a descriptive, language intensive activity. The costs of development and perhaps more importantly maintenance of linguistic resources needed to localise clinical systems are high. Any practical approach to the management and exploitation of linguistic resources in large-scale clinical information systems must be based on common methods and internal representations for linguistic information. This information must be reusable across a wide range of systems and local variants of those systems, and the cost of maintaining that information must be separable from those of maintaining the rest of the system.

Significant regional differences in linguistic usage exist even within single languages, and even more so when minority languages are taken into consideration. Not only are we faced with differences in language, but also there are differences in culture. Though the Dutch and Flemish world share most of their words, the actual daily meaning of these words may be markedly different! To be truly successful, a research programme of linguistic engineering in medicine must have a strategy for recognising and managing regional as well as national linguistic differences. These considerations further complicate and increase the expense of 'localisation' of products for the European market.

There is an increasing awareness of the crucial role of a 'common medical language' for the further development of medical information systems. The time of 'black box thinking': 'Natural Language Processing will solve the problem' is over. Moreover, NLP also needs strong and large domain models in order to be effective and reliable. Introduction of knowledge-based systems is severely hampered by the lack of common terminology.

Traditionally classification systems have been developed for a number of quite different purposes: Statistics (ICD, etc.), Nursing (ICNP, NANDA, etc.), Human Functioning (ICF), Case Mix (DRG, etc.), Procedures/Reimbursement (ICPM, CADM, CPT etc.), Epidemiology (ICPC, etc.), Literature retrieval (UMLS), Quality of Care, Protocols, Medical Records (Read, SNOMED, Gabrieli). Though from this list it is clear that there are standards for some elements of the Medical Record, specifically elements that characterise the patient, such as signs and symptoms, are hardly standardised.

If terminology is to be used in automated systems, terminology development can certainly not be seen in isolation:

- First of all it requires insight in the business process of the health care organisation. How does information flow through the organisation, and between organisations? So we have to address the problem of what is generally put under the umbrella of Messaging.

- Secondly for documentation we need storage structures. As stated in the introductory lines, there is a lot of activity at present in the area of structuring the Electronic Medical Record. There clearly is a relation between terminology and the names of structures in the record. There is on the one hand a number of actors in the field of record structures (HL7, OpenEHR), on the other hand a diversity of developers of terminologies. There remains a grey area, where it is not clear who should have the lead.

- Thirdly we need the terms and a delivery mechanism.

3.1 Some current bottle necks in terminology

- Specialist groups develop code lists for (national) registries; these code lists are often incommensurable with (internationally) accepted classifications;

- Classifications are frequently developed for other purposes than registration of the primary care process;

- Specific needs within certain sectors or specialties are not met by current classifications. This leads to proprietary lists and additions to existing coding systems;

- Proper methods for controlling extensions to classification systems are missing. These methods are needed for comparability, aggregation and data exchange;

- Different sectors use different standards for reporting of the same rubrics (e.g. history, diagnoses)

- Existing classifications are often inappropriate for information exchange between sectors because there is no mapping between the classifications used;

- Existing classifications are often inappropriate for registration of the care process, as they lack a sufficient level of detail;

- Existing classifications are often inappropriate for registration of the care process, as they have not been developed specifically for a given sector (ICIDH, …)

- Proper registration of the context of care requires a suitable record structure. This structure is missing. This topic must be thoroughly addressed when an EPR is being

developed in the Netherlands;

- Coding of similar information does not necessarily lead to similar results. For unequivocal registration it is mandatory that similar data be coded in the same way. Support is needed to realize that.

3.2 Registration

Terminological systems do not support clinicians in determining the condition of a patient or the medication to be prescribed.

These functions could be provided by Decision Support Systems, which in turn can use a terminological system.

A Terminology Server may help in naming data, when a clinician knows the actual situation, a Terminology Server can be used to find the relevant terms. If, for example, a clinician diagnoses a viral infection in the liver, then a Terminology server can provide the name "viral hepatitis", if this term is present in the Terminology.

It is possible that a Data Entry System is being used to request for further specification, for example to ask for a specification of the virus causing hepatitis. In this case, the Data Entry System will query the Terminology Server for attributes of hepatitis that can be specified (e.g. aetiology), and will retrieve a list of relevant attribute-values (e.g. viruses that can cause hepatitis). The clinician then needs to determine whether this information is known and relevant given the situation, as determining aetiology or specific pathology, can be time consuming, costly and/or aggravating to the patient. This information will be collected only when it is required to determine appropriate treatment, or is otherwise in the interest of the patient. No system, except for a Decision Support System eventually, supports determining the relevance of data for clinical practice in a given context.

As a clinician may or may not be able to describe a patient in detail, it is required that information can be specified at the relevant and available levels of detail. A high level of detail for various attributes may lead to a large number of potential values (e.g. for diagnoses, procedure), which requires the possibility of constructing such highly specific detailed values, by combining basic elements. This is called post-coordination.
In the example of viral hepatitis, this could be combined with course (acute, chronic, acute on chronic), severity (mild, severe, …), aetiology (a specific virus), and so on. Existing terminologies can suffer from any of the following (related) problems:

- Items can not be recorded at the required level of detail

- There is a number of concepts in a terminological system that are equally appropriate in describing a patient. This occurs for example in CPT(Current Procedural Terminology), where a coding guideline describes which code to use in a given situation. These guidelines however change over time, resulting in different coding of the same surgical procedure. This leads to loss of continuity, as a seeming change in treatment is actually a change in coding practice. Other coding systems may fully lack coding guidelines, leaving it to the individual clinician to select a proper code.

Differences in individual or regional/national preferences lead to incommensurability of data.

- A concept that matches the actual situation has too much detail. E.g. a terminological system contains the concepts *infectious hepatitis*, *acute viral hepatitis*, *chronic viral hepatitis*. In this case, *acute infectious hepatitis* cannot be recorded; instead a too generic or a too specific concept needs to be chosen. The problem is in the granularity.

- If a system supports post-coordination, it is important that post-coordinate concepts with the same meaning are detected to be equivalent (see 7.1.1.2). It is perhaps even more important to be able to distinguish post-coordinate concepts that seem to have the same meaning, but that are actually different. Suppose enteritis (Inflammation of the intestine) is defined as an inflammation, localized in the intestine. A clinician could register an inflammation of the duodenum as "inflammation localized in the intestine", if the exact location is not (yet) determined. The latter is however not an enteritis (it involves only an undefined part of the intestine), hence it must be possible to distinguish these. However, a clinician could register enteritis in the same way as he has registered the duodenitis. The issue here is about expressive power of the terminology.

A Data Entry System that stimulates registration of patient information at the proper level of detail will result in registered data that has added value both for daily practice and for research because it is more precise. This will in turn benefit registration of the information.

Generally agreement exists that reasons for registration must be clear and valid to the people that register information. Quality of the registered data depends not only on the underlying terminological system and services, but even more on the participation of clinicians to register information conscientiously. Providing feedback (e.g. results based on the registered data) is important to stimulate clinicians.

3.3 Communication

The enormous size of contemporary terminological systems can hinder the proper use of them. Clinicians must know the required level of detail in a certain context. Too much detail may require too much effort of a clinician, whereas too little detail hampers the usefulness of data. Communication between specialties requires knowing a priori the required level of detail for the recipient of the information, e.g. an extensively described operative procedure is not relevant for a general practitioner who is only interested in the outcome of the procedure.

3.4 Aggregation

A good terminological system provides the possibility of grouping registered data in a variety of ways. It is however important to realize that:

- The registered data are not necessarily complete. The number of patients for which a certain disorder is registered may be smaller than the number of patients actually

suffering from the disorder. Patients can be completely missing in a registration, or the registration of a disorder seemed irrelevant earlier.

- There may be a variety of ways to register data. Fever can be recorded as a finding, e.g. based on a terminological system. It can also be that a patient's temperature is recorded that indicates fever (e.g. 39 Centigrade). Another possibility is that a patient gets medication (e.g. antipyretics) from which can be deducted that a patient (possibly) has fever.

- Generally, not mentioning a disease or symptom is not equivalent to the absence of it. Hence, "no mention of fever" is different from "no fever", which might be explicitly registered.

3.5 Examples

Below we give three examples demonstrating problems with medical language, and rules that apply to the usage of that language.

3.5.1 Expert agreement in Current Procedural Terminology evaluation and management coding [63]

BACKGROUND: Available data suggest that physicians are accurate in approximately 55% of Current Procedural Terminology (CPT) evaluation and management (E/M) coding for their services. This accuracy is relative to observers' or auditors' assigned codes for these services, a group that has not been studied for their consistency in application of the CPT E/M coding guidelines. The purpose of this study was to determine the level of agreement of certified coding specialists in their application of CPT E/M coding guidelines.

METHODS: Three hundred certified professional coding specialists randomly selected from the active membership of the American Health Information Management Association were sent 6 hypothetical progress notes of office visits along with a demographic survey. The study group assigned CPT E/M codes to each of the progress notes and completed the demographic survey.

RESULTS: Coding specialists agreed on the CPT E/M codes for 57% of these 6 cases. The level of agreement for the individual cases ranged from 50% to 71%. Relative to the most common or consensus code, under coding of established patients occurred more commonly than over coding. In contrast, for new patient progress notes, over coding relative to the consensus code was more common than under coding. CONCLUSIONS: There is substantial disagreement among coding specialists in application of the CPT E/M coding guidelines. The results of this study are similar to results of prior studies assessing physician coding accuracy, suggesting that the CPT coding guidelines are too complex and subjective to be applied consistently by coding specialists or physicians.

3.5.2 Diversity in death certification: a case vignette approach.[64]

Given the same case history information, physicians showed great variation in wording and diagnostic semantics in death certification. The rates of correct certification format and concordance with referent UCOD varied with the level of complexity of the causal sequence of death. The greatest source of diversity was choosing between an acute

condition of a chronic disease and the chronic disease itself, and between competing prominent co-morbidities.

3.5.3 The Challenge: Delivering Clinical Decision Support for Antibiotic Treatments Despite Disparate Systems

The timely and accurate use of antibiotics is one of the most pressing issues facing the healthcare industry today. According to a June 2002 article in the Annals of Internal Medicine, 75 percent of hospitalized patients receive antibiotics and 50 percent of that use is inappropriate. A study published in the Journal of the American Medical Association has estimated that between 20-30 percent of all adverse drug events are antibiotic-related. The misuse of antibiotics can impact a patient's health and length of stay and can also contribute to increasingly worrisome public health issues such as the growing problem of anti-microbial resistance.

More than 15 years ago, the founders of TheraDoc began to conduct pioneering work in the field of antibiotic-related decision support, creating software that provided healthcare professionals with knowledge-enriched, disease-specific recommendations for antibiotic treatments, tests and referrals based on individual patient profiles. A January 1998 study in the New England Journal of Medicine that examined the implementation of this software in one hospital found significant benefits, including:

- 85 percent reduction in adverse drug events
- 33 percent reduction in mortality
- 7-day reduction in length of stay
- $18,000 reduction in cost of hospitalization

Founded in 1999, one of TheraDoc's greatest challenges was to create expert system technologies that could be easily transferred to other institutions. To deliver patient care benefits, the software had to be able to integrate, standardize, and create actionable data from various systems in numerous departments, including admissions, pharmacy, laboratory, radiology and microbiology.

The solution to that challenge was facilitated with the integration of SNOMED CT. With SNOMED CT's ability to normalize data from various sources, TheraDoc's decision support software can be replicated quickly and cost-effectively without building a complex new system from the ground up in every single hospital.

4 Requirements

Terminology is a pervasive issue in healthcare, it lives everywhere, in the contents of textbooks, scientific journals, databases, messages, health cards, clinical guidelines, elements on screens of clinical information systems and so on. The profession is not always aware of the importance to have unity in language. Unity does not mean we (or our systems) all have to speak exactly the same language. What we do need is unequivocal transcription from one place to another without information loss, or if there is information loss, we need to know which information is lost.

Healthcare Information Systems are moving rapidly from management and administrative functions towards support for clinical care, quality assurance, and resource management. The market for administrative hospital systems is rapidly becoming saturated whereas the market for clinical systems has just begun to grow. The movement towards clinical systems is motivated by at least four key trends:

1. The move towards 'evidence based care'.
2. Pressure for greater cost effectiveness and more coherent planning in health care.
3. The shift from 'doctor centred care' to 'patient centred care', which requires increasing collaboration amongst different professionals and transfer of information between individuals, wards, institutions and sites.
4. The trend away from monolithic hospital wide single-vendor solutions towards heterogeneous multi-vendor solutions.

All of these changes depend on:

- Effective capture of detailed clinical information at the point of care;
- Effective communication of clinical information between users with different needs, different perspectives, and different working practices;
- Effective communication of clinical information between heterogeneous applications.

Many studies stress the fact that capture of clinical information at the point of care is essential for new technologies to be successful. At the same time it is still very awkward for care providers to review and enter substantial amounts of data during a patient encounter.

4.1 Who are the customers

This study looks at the need for a detailed clinical terminology from a broader perspective than just those of the now eminent health record. In the context of a record one might argue that notions like 'family history' do not need to be represented in a clinical terminology, because that is taken care of by mechanisms like headers, archetypes, or some other record structure. In document retrieval the notion of 'family history of diabetes' is quite relevant. Other users do need standardised data too. National registries are notoriously bad with respect to quality. The national offices of health statistics, as well as international organisations like Eurostat and OECD depend on these data. Quality data is necessary for direct patient care, but also important for scientific research (epidemiology, medical technology assessment, etc). Another still

underdeveloped use of detailed recording is the internal benchmarking and quality control of the healthcare delivery. Pressures from patients and insurance companies will lead to more detailed recording. The healthcare provider will rightfully state that his role is to cure or provide care for patients rather than being a clerk. The major challenge is to provide an environment and incentives to the care provider that allows for quality routine recording. It is obvious that we should not only look at clinical terminology in the context of ICT implementations, but also be aware of secondary uses of healthcare data.

4.2 Editorial Control

4.2.1 Influence of professional bodies

It is obvious that the terminology should match the needs of the healthcare provider. If terms are clearly recognised as part of their daily professional communication, it is more likely that the terms will be used properly. They should feel a certain ownership of that terminology. Ideally the national professional societies should have a strong position in the editorial boards. For commercial offerings, which often are black boxes, such an influence does not exist. In the case of large international suppliers the potential influence of a small country like The Netherlands is likely to be rather limited[1]. In this context it is worth noting that there is a great fear of market dominance by the large USA suppliers of terminologies and related software products. The parallel with Microsoft/Intel is often mentioned.

4.2.2 Maintenance/Training

Licensing a terminology from a third party reduces the efforts of maintenance, although a central point that collects and evaluates requests for changes, and communicates these with the terminology provider, is still necessary (see 4.2.3). It will also create a dependency from this third party, and decrease flexibility. On the other hand, maintaining a proprietary terminology will be a huge task, with corresponding costs, and less possibility for international comparison and cooperation.

In traditional classifications the update frequency of procedure classifications is markedly higher than in diagnostic classifications. For clinical terminologies the same is true, typical update frequencies are in the order of twice a year. Because applications and clinical terminology are becoming ever more intertwined, version control becomes a serious issue. It is not only updating of some lists, but terms to describe record elements, protocols, screen items to name a few might be subject to change. Before introducing a clinical terminology the change management needs to be organised and be in place!

Because of the pervasive nature of a clinical terminology many need to be trained, care providers, software vendors, and other users of health data. The amount of training, and hence the costs are dependent on how complex the terminology is delivered.

4.2.3 Expertise Centre

The sheer size of a clinical terminology, and the fact that its intended use is across sectors calls for a national reference point for both content and implementation issues.

[1] As an early adopter there might be opportunities for increased influence

The necessity of frequent localization demands short communication lines. Even in this electronic era it is not advisable to depend solely on a supplier abroad. In order to cover sufficient depth one should think in the order of 3-4 persons staff for such a function.

4.3 Economics

Introduction of a large clinical terminology bears considerable costs, not only for licenses and implementation but also for a large extent due to the necessary organizational changes. Some care providers are likely to get more administrative duties. On the other hand there are benefits in the immediate availability of data recorded elsewhere. It is expected this new approach will lead to substantial change in the health care delivery process. In reality this is a two way process. Necessary changes in the healthcare delivery process asks for larger formal terminologies, on the other hand larger terminologies demand other structures to properly cope with complexity. If the record is just an electronic copy of the paper record, there is hardly added value to be expected. Record architectures will need to be changed to achieve the full advantage of new terminologies. If a formal structure is added, proper terminology becomes a necessity. Expected benefits are in improvements of care process, improved comparability, better data for research and epidemiology, and in safety of patient care. This all will cost real money. Care providing organizations must be convinced of the added value before they will consider implementation.

Estimates have been made about the costs of developing and maintaining a large clinical terminology, figures range in the order of $20M to $35M[2], where maintenance costs are in the order of 5-10% of the total. If the terminology is not developed locally, costs of translation and localization are estimated to be in the order of 10-20% of development costs. There are different options for licensing, either each institution, or vendor acquires a license for usage, or alternatively one can opt for a national license. Countries with socialized health care, like UK and Australia, tend to opt for national licenses. National licenses may lead to cost savings of over 75%, this seems attractive from a financial point of view. It remains to be seen how such a model fits with the national free competition rules and those of the European Union.

New terminologies are delivered via so-called terminology services. Next to the expense of developing and maintaining a terminology, also similar expenditure is necessary for terminology servers. There is no single correct solution for implementing terminologies. As stated before the new large terminologies will be made available via Terminology Servers (TeS). Apart from the specification of the Lexical Query Server (LQS) from the Object Management Group (OMG), there are no standards for TeS implementations. And even LQS only serves the very simple first generation terminologies. To reduce costs of implementation it is advisable to aim for at least a national specification for a TeS interface.

The figures presented here are only indicative. In comparison to present spending on terminology[3] the amounts are huge, in comparison to the total healthcare budget it is very small. If considered as part of healthcare ICT expenditure, it still is very small. If we

[2] Note: SNOMED CT merges 2 precursor works each of which required that level of investment.

[3] It must be noted that traditional classification development is to a large extent the work of volunteers. This bears substantial hidden costs.

do compare healthcare ICT expenditure with other sectors it is minimal. If the use of a reminder system based on a proper terminology (see example 3.5.3) would help to reduce costs[4] of errors in medication by only 1%, quite a bit of costs for that terminology would be recovered. Not to mention preventing enormous personal grief of patients. In the Dutch healthcare system there are many different players and payers. Some players will benefit, others have to pay in whatever form. For a successful implementation agreement of all parties concerned is imperative.

4.4 Deployment

It is current practice in industry that the supplier takes initiative in setting terms and conditions under which a certain product or service is licensed. Specifically in the software industry where one supplier delivers to many customers these terms and conditions are mostly one-sided in favour of the supplier. In the situation where substantial purchases are made, the game is different. The purchaser defines the product in terms of specifications. Quite often such a purchase goes for public tender. The end result will be a more balanced agreement of rights and obligations for both the purchaser and supplier.

In case of a large national investment on terminology, be it by purchasing off-the-shelf, or by commissioning a new development, the 'tendering model' seems most appropriate. It is likely that the sheer size of the planned investment even does prescribe public tendering because of EU regulations. In the context of tendering there are a number of requirements to consider ensuring a high quality stable product:

1. ***Safety of patient data*** The first and most important requirement is safety of patient data. The discussion on patient data safety is predominantly focussed on issues like confidentiality. The quality of data, and the quality of derivatives not only matters for laboratory data, but also are equally important for terminology. The product should provide sufficiently precise terms for recording. The exact meaning should be obvious from the term itself and/or the context in which it is presented. There is a second aspect of safety specific for third generation terminologies. Those packaged in software (Terminology Server or TeS) do have a defined functionality. Such systems can make statements about terms it contains. If e.g. a TeS drives a reminder system, its answers should be complete and correct. Though formal proof is impossible, purchasers and suppliers should specify metrics for assessment of quality.

2. ***Continuity of access to patient data*** For pragmatic reasons terms are mostly stored in coded form in the patient record. For future reference of those records the terms and its links to codes are necessary. If for some reason the agreement between supplier and purchaser stops, continued use of the terminology for retrieving existing data should be secured. Also the right to convert existing data into some other terminology should be established.

3. ***Multiple suppliers*** In our open market terminology should be available from multiple suppliers. This applies both to the software in which it is delivered as well as

[4] Estimated by NICTIZ at €300M annually for The Netherlands

to the term sets used. It must be possible to use a term set from supplier A next to a term set of supplier B within one software environment.

4. ***Specification and Certification*** The terminology must be explicitly specified in terms of its objectives and how it should be implemented. The supplier must also deliver the criteria for certification of the product, and rules to define certification data. The supplier may specify additional rules and tools for safe third party modifications/extensions on his terminology. As a matter of principle a trusted external party conducts the certification. A supplier will only be held responsible as far as his product has been certified.

5. ***Ownership*** Ownership of a terminology should only apply to the delivered/updated term collection as a whole, and the way in which it is organised. Individual words and phrases are born from the medical community, and as such are part of the public domain in the same sense that there is no owner of the English language.

6. ***Extensions and updates*** Healthcare language is highly dynamic, often with needs for local variants due to different healthcare delivery systems. Purchaser must secure that he has both the rights and the tools to make corrections and extensions. Purchaser should be encouraged to make corrections and extensions available to the supplier. Purchaser should retain the right to re-use extensions elsewhere.

7. ***Usage*** The terminology should be available for all healthcare related work and for research. This does include alternative usage in different kinds of software. Provisions must be made to avoid infringements on rights of both the supplier as well as the parties on whose behalf the purchaser enters into the agreement.

8. ***Availability of resource*** Both supplier and purchaser have an interest in safeguarding their investment. License agreements are likely to be made for a longer period. The time window for a purchaser is determined mainly by the time it is expected to take to change to a different supplier. In the early stages of implementing the terminology this means a shorter contract period. Later if the business process cannot be disturbed the purchaser needs a further horizon in the contract. The supplier needs to survive, and hence is interested in continued return on investments. Availability is also the purchaser's interest.

The aforementioned requirements strongly suggest aiming at some kind of open source licence agreement. Open source will most definitely not be free of cost; a complex terminology is required, and we should not expect this to appear without significant investment!

It will not be an easy task to meet all the requirements mentioned. There are tensions between the demands of multiple suppliers and safety of patient data. Certification and multiple suppliers is also a difficult issue. The model of a single market-dominating supplier may seem attractive; the situation in the software marketplace has shown that this does not necessarily guarantee quality.

Any national license entered into for the use of SNOMED CT or other terminology should include appropriate protections of the integrity of the terminology, perpetual use rights,

potential for a competitive supplier market, documentation and implementation specifications, as well as appropriate opportunities for development of extensions to meet local needs.

5 Technical aspects

The terminology the user sees in applications is not always the terminology as bought off the shelf, but rather a collection of terms familiar to a specialty provided by the software vendor. These terms are presented as pick lists, or as labels on computer screens. These lists and labels are reflected in the databases. This collection of terms usually is referred to as *interface terminology*. When the interface terminology is mapped to a *reference terminology* (e.g. OpenGALEN, SNOMED CT), it becomes possible to exchange data between dissimilar products. The reference terminology takes care of the proper definition and organisation of terms. The distinction between reference terminology and interface terminology is somewhat artificial; most suppliers deliver both in one package. However, the distinction can still be made on the level of tasks, where the central task is to develop and maintain the global reference terminology and the local task to make correct mappings from a software product to the reference terminology. In such a situation it must be possible to make a new concept on the fly in the reference terminology if it is not there yet. The reference terminology has become invisible for the user, and has become middleware.

There is increasing agreement on which aspects of terminologies are important. Many of these aspects relate to the *formalism*(s) for representing concepts, terms and cross-mappings to other coding systems. Terminologies are required to be concept-oriented, meaning that terms are explicitly separated from the concepts they represent. This enables representing concepts with multiple synonymous terms, possibly in multiple languages. The formalism determines the possibilities for representing *knowledge*, and contributes to the *functionality* that can be provided by the terminology. We will therefore focus on formalisms, and the required capabilities.

5.1 Description logics

An important formalism for representing concepts and relationships between concepts is known as *description logics*, which refers to a family of knowledge representation languages that give a formal basis to traditions like frame-based systems, semantic networks, KL-ONE like languages. It is closely related to propositional modal logics. One of the main features of a description logic is that it has well-defined semantics, the other is that it provides a basis for inference services, e.g. automatic classification of concepts.

However, the use of description logic alone does not guarantee that the resulting terminology is correct and complete. There does not exist (nor will there exist in the foreseeable future) a method to prove that the contents of a terminology are correct and complete. As with any other formalism, the results of the modelling process depend on the proper modelling of concepts. Inferred classification is hampered when so-called primitive concepts are abundantly used, or if concept definitions are not modelled in a consequent way. The process of knowledge modelling is crucial for exploiting and controlling the power of a description logic.

5.2 Messaging

The usage of a terminology depends on the ways in which multiple applications can interchange terminology-based information. This means that messaging standards must be able to deal with the codes and terms used in a terminology, and that applications agree on the used terminology. Another option is the use of an intermediate (reference) terminology that enables "translation" of codes from one system to another.

The emergent de-facto standard for health communication is HL7. In the current version (2.4) of HL7, the messages consist of segments with fixed-length fields. The maximum length of the diagnosis code in the DG1 (Diagnosis Information) segment is limited to 60 characters, and the diagnosis description to 40 characters, which may occasionally be too short. The messages in the next version of HL7 (version 3) will be based on XML, which will remove the need for limits on the maximum length of any field.

5.3 Electronic Patient Records

To fully exploit the strengths of advanced terminologies, coding should take place during the care process itself and by the treating physicians. This is a fundamental change in the coding practice of today, where coding is performed mostly post hoc by coding experts. This requires that physicians are trained in coding, but also dictates very intelligent and user friendly software that guides physicians through the overwhelming contents of the terminology.

5.3.1 Storage of coded information

EPRs must provide the possibility of storing terminology-based data. The current generation of coding systems mainly use codes that are limited in length and format; e.g. ICD-10 codes start with a capital letter, followed by two digits, and possibly a period and another digit. Using a terminology that allows for post-coordination (7.1.1.2) may result in codes of any length because of using code sets. Besides, it is advisable to also store a textual representation in the EPR, in order to ensure permanence and availability of the registered patient data, but also for legal issues. In some countries it is required by law that textual strings as seen by the care provider are stored in the record.

5.3.2 Integration of terminology browser

Terminology based registration of patient information in everyday clinical care requires a terminology to be integrated into the care process, in other words the terminology must be available from within other applications, like the EPR data entry system. Functionality must be provided either as an application specific extension, or as a terminology specific module being incorporated into an application. The former scenario requires EPR vendors to provide access to the terminology; the latter requires integrating an independent terminology browser.

5.4 Training

As mentioned in the previous section, introducing terminology-based registration of patient data involves a major change in registration practice. Instead of post-hoc coding of free-text data, data is entered in a structured manner, being coded directly. This will

involve training of users on the preferred ways of performing terminology-based registration.

To our knowledge, no structural evaluations of introduction of terminology-based registration have been performed as yet. Important issues in such an evaluation are the increased effort for clinicians to register data, and the improvement of quality of data.

6 What is happening outside The Netherlands

Most developments in clinical terminologies are based on perceived needs. Quite often developments have started around a single health care domain (e.g. pathology) or by visionary individuals who want to catch a great level of detail in their recordings. Large scale terminologies like Read, SNOMED, and OpenGALEN, don't seem to have a thorough business case analysis to motivate their development. As a matter of fact the Clinical Terms/SNOMED merger was motivated on the UK side primarily by the expectation to get a better grip on terminology development, and hence on its costs.

After the finalisation of this inventory it was brought to our attention that some in house experimentation with SNOMED had started in The Netherlands. In the context of this report the ePath-poject at the University Medical Centre Utrecht is worth mentioning. An event-based registration system for traumatology is under development there. This application area is very data-intensive, and does need a broad coverage of terminology both in scope and detail. Because of the nature of a traumatology department, there is a need for communication with other departments to which the patient is transferred eventually. Also communication to primary care is an important link for the department. This makes traumatology an interesting candidate area for further pilot studies.

6.1 SNOMED

SNOMED has an international position. The website mentions 36 countries in which any version of SNOMED is being used somewhere: Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, China, England, Finland, France, Germany, Greece, Hungary, Iran, Ireland, Israel, Japan, Malaysia, Netherlands, New Zealand, Norway, Poland, Portugal, Russia, Scotland, Scotland, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Turkey, and Wales.

Currently, SNOMED focuses on languages that cover large populations, or that are of relevance for other reasons: Spanish, German, and Dutch (among others). On the other hand, some countries are actively involved in terminology issues, considering using SNOMED as their national terminology. Some of these have been involved in the SNOMED CT Alpha evaluation: UK (22 sites), Australia (9 sites), US (8 sites), Iceland, Germany, and Netherlands.

Today, SNOMED CT is licensed and used for a variety of clinical applications that include Physician Order Entry for laboratory, imaging and drugs, public health reporting, telemedicine, genetic databases, therapeutic decision support, cancer registry reporting, emergency room charting, 3D imaging and auto-indexing, electronic health records, problem lists and disease reporting.

More specifically, as advertised since its release in January 2002, the College concluded over 50 supplier licenses for SNOMED CT. The state of incorporation into vendor systems varies and is largely dependent on the vendor's development cycle.

In this section we will have a look at a number of those countries and the activities they have undertaken.

6.1.1 Australia

The Australian government is currently investigating if and how to introduce a national clinical terminology. They see two options:

1. do nothing. Governments have had not much success in picking a winner. The idea is that through competition in an open market the 'best' terminology will emerge. At the moment, SNOMED-CT seems to be the only candidate, but it has not yet been demonstrated in widespread use and it is unclear if it will eventually be the winner. In the end, only the users, i.e. the clinicians, themselves can judge how appropriate a terminology is.
   However, this option may lead to several incompatible 'reference' terminologies, which will make it very difficult, if not impossible, to aggregate data. It will also bring great costs to the government in mapping the various terminologies to each other, which is vital to epidemiological studies.

2. mandate a standard reference terminology for Australia. There is considerably agreement on the need for a single standard reference terminology, preferably in the public domain. This removes the issue of a standard terminology from the market, and is likely to bring considerable external benefits. A number of other projects depend on the existence of a standard terminology, e.g. Health*Connect*[5]. Within this option there are three scenarios (order is based on the costs, where scenario a. has the highest costs):
   a. acquire SNOMED CT
   b. develop a reference terminology analogous to SNOMED CT, but appropriate for the Australian needs
   c. develop a minimal reference terminology based on work in progress on the basis of ICD-10-AM
   d. Based on past history, where Australia started with a 3M DRG system and ended up specifying its own AR-DRG system, it may be possible that Australia will develop its own terminology.

6.1.2 United Kingdom

In the UK, the NHS is involved in the merging of CTV3 (Clinical Terms version 3, formerly known as the READ codes) and SNOMED RT. The result, SNOMED CT, is the preferred clinical terminology for the NHS.

"SNOMED Clinical Terms creates a single unified terminology to underpin the development of the integrated electronic patient record by providing an essential building block of a common computerised language for use across the world. The NHSIA and CAP both have current terminologies, which have been recognised as world leaders in the field. The merging of the two works will create the most comprehensive language of health to support the computerised patient record."

Founding Principles [http://www.nhsia.nhs.uk/SNOMED/pages/principles.asp]

- Improve patient health and medical science through the availability throughout

[5] HealthConnect is the proposed national health information network to facilitate the safe collection, storage and exchange of consumer health information between authorised health care providers.

healthcare of a consistent terminology for data capture, retrieval and analysis, to support clinicians in providing good clinical care and other secondary purposes.

- Promote the adoption of the SNOMED Clinical Terms as the premier terminology for healthcare and veterinary practice throughout the world that can be adapted for national purposes if necessary.

- Seek the continued guidance and incorporation of knowledge and experience of the clinical specialities through active involvement on the Editorial Board.

- Enable the scientific development of SNOMED Clinical terms while attempting to recover costs through effective licensing strategies and market education.

- Respond in a pro-active way to user needs with the needs of the USA and UK taking precedence equally over other markets in the world.

- Collaborate with suppliers (in UK as well USA) to minimise excess development costs (and thereby user costs) in upgrading to SNOMED Clinical Terms.

- Ensure that the breadth and comprehensiveness of SNOMED Clinical Terms will be greater than current versions of Read Codes and SNOMED.

- Describe and follow clearly defined processes for the integration of user feedback and provision of updates.

- Work with users and system suppliers to develop a sound migration strategy that will clearly define the upgrade path from the current versions of Read codes and SNOMED to SNOMED Clinical Terms.

- Periodically review licensing and pricing strategies as adoption of SNOMED Clinical Terms expands across the globe, in order to reduce barriers to use while supporting ongoing development and growth

The effort involved in the merging process and the introduction of SNOMED CT is large. A variety of activities has been deployed, and a number of project teams has been formed. These include (but are not limited to):

- Performing a "Controlled Health Terminology Survey" [C. Price, 2001]
- Development of "An Evaluation Method for Clinical Terminologies – Technical Component" [NHS, may 2000]
- Start up of a project on "Specifying Clinical Term Sets for Information Requirements (SPECTRE)" [NHS, January 2003]

Efforts can be grouped under the following headings:

Merging CTV3 and SNOMED RT. This has been realized, but work is ongoing on validation of the completeness and correctness of the merging.

Development of subsets. The subset mechanism of SNOMED CT makes it possible to define specialty-specific subsets of SNOMED CT concepts and terms. For a number of specialties (cancer, diabetes, mental health, coronary heart disease), such subsets are being developed.

Maintenance. Whereas not only medical knowledge is constantly evolving, but also medical practice may change over time, continuous maintenance of the clinical content of SNOMED CT is required. Although the SNOMED organization performs such monitoring, they are also depending on input from partners

Implementation. As NHS will eventually mandate registration of patient information based on SNOMED CT, implementation of SNOMED CT needs to be realized. This implementation will not be performed by NHS, but by EPR vendors.

It needs to be noticed that currently there are no differences in usage of the SNOMED-coded data as compared to the CTV3 coded data. Though it is expected that the use of SNOMED CT will create a range of possibilities w.r.t. research and evaluation of care, there are as yet no initiatives to explore these.

6.1.3 Sweden

In Sweden a project called REFTERM at Karolinska Institute has recently started. The REFTERM project aims to evaluate the possibilities for establishment of a Swedish reference database for medical terminology to facilitate health communication, research and follow-up based on the international work with SNOMED from the College of American Pathologists. The project at Karolinska Institute I is conducted in co-operation with the Stockholm Country Council (the regional health authority) with the first objective of building a base of knowledge on Snomed and to perform a limited evaluation of its usability in a few contexts where the basic IT-solutions for communication are already present and are developed in other project forms. Collaboration will be established with other central projects for information exchange of the region such as Julius and a Common Care Database on a Middleware platform. The REFTERM project will also evaluate the feasibility of using the SNOMED system for research, primarily clinical research but also certain other basic biomedical research, e.g. in the borderland between genetic knowledge and clinical information (functional genomics). Co-operation with other Swedish actors in the Terminology field is established through a reference group with participation from Universities of Linköping and Gothenburg, Swedish Society of Medicine, National Board of Health and Welfare and Terminology Centre (TNC). Together with Linköpings University we also have an exchange of experiences and joint development of tools to handle SNOMED.

The project has not yet published any results. The Karolinska Institute has expressed interest in further European cooperation for evaluation SNOMED CT.

6.1.4 USA

In this project we planned at the end of March 2003 a one day meeting in The Netherlands with some representatives from SNOMED and representatives of Health Language, a supplier of terminology services. Theme of the planned discussion was the potential added value of SNOMED CT in cross sector communication. Due to the geo-

political situation the meeting was cancelled. However, much has transpired since that time.

6.1.4.1 U.S. National Library of Medicine (NLM)

On July 1, 2003, Secretary of the U.S. Department of Health and Human Services, Tommy Thompson, announced that his agency signed a licensing agreement with the College of American Pathologists for SNOMED CT. In his speech at the U.S. National Healthcare Information Infrastructure Conference, he proclaimed that "this wonderful groundbreaking agreement will improve the quality of care by allowing electronic communication among all providers." Most of the issues raised in section 4.2 and 4.4 of this report can be addressed by the details and terms of the agreement. The agreement can be found on www.nlm.nih.gov and www.snomed.org.

Following the announcement, the Gartner Group, a major I.T. industry analyst, issued a report on July 3, 2003 entitled "Use of SNOMED CT Should Improve Healthcare Information" stating "this step will vastly improve the exchange of clinical information." They further go on to say that "this announcement guarantees that SNOMED CT will become the defacto standard for representing medical concepts . . ."

We did interview a representative of Kaiser Permanente, a large US HMO (Health Maintenance Organization) that has actively participated in the development of SNOMED and provides health care coverage to 3% of the U.S. population. Topic of the interview was to get some insight in potential added value of SNOMED CT. How much of the savings can be attributed to SNOMED only could not yet not be said. At present Kaiser is converting to a new health record system from Epic Systems Corp. The conversion is expected to be completed by 2005. At that time Kaiser expects to be able to be more explicit on the added value of SNOMED CT on clinical outcome.

6.1.4.2 Kaiser Permanente Group, Colorado[6]

In 1998 after ten years of design, development and implementation, Kaiser Permanente, in the state of Colorado, started using a Clinical Information System. The project is praised as a successful state of the art example, both functionally and technically, of an electronic, clinical system for the exchange, storage and management of information. The system provides all care-providers in the state of Colorado interactive and on-line access, including updating, to clinical information. The kernel of the system consists of a large clinical database, diagnostics and support via SNOMED[7], several clinical applications, an integrated vocabulary that generates standard terminology from plain text reports, order communication via HL7, a terminology server that translates diagnostic codes to financial and procedural codes. Approximately 500 physicians and 2000 supporting personnel use the system, while caring for 350.000 patients (one patient file every second). Evaluation analyses have shown that the system saves the organization 8-10% exploitation costs, of which is 4-15% direct savings on physicians and research and treatment departments, and almost 60% on administration.

[6] This is a summary of a site visit by HL7-NL

[7] We were unable to verify which version of SNOMED has been used. We assume this is SNOMED RT, a smaller and simpler version prior to SNOMED CT.

6.2 Other Systems

The major focus of this study is on SNOMED CT. For reference purposes we have collected public descriptions of a number of commercial and OpenSource offerings. Those descriptions are available in Appendix C: Other systems.

We have added/requested from vendors some screenshots of their systems focussing on heart, to give a first impression on the kind of knowledge in each resource. Apart from the open source systems is hard to get in-depth insight of structure, functioning and content of the separate offerings. A standardised method for describing and assessing terminological systems is still lacking. Notions like 'semantic network' and 'description logic' are used, but these do mean quite different things between vendors!

This appendix describes a number of alternative terminologies. The information in this document represents (by default) a snapshot in time. For the most recent information the listed web-sites should be visited.

It should be noted here that from the systems below LinkBase has cross-links to the Dutch language. All the 80.000 index terms of the Dutch Journal of Medicine (NTvG) have been cross-mapped to the MESH index via UMLS.

Terminology	**web-site**	**distribution**
ID MACS	www.id-berlin.de	commercial
LinKBase	www.landcglobal.com	commercial
LOINC	www.loinc.org	open-source, free of charge
MED	www.cpmc.columbia.edu/homepages/gum7001/topics/index.html	Free ?
MEDCIN	www.medicomp.com	commercial
MESH	www.nlm.nih.gov/mesh/meshhome.html	open-source, free of charge, under MeSH Memorandum of Understanding
OpenGALEN	www.opengalen.org	open-source, free of charge, under GALEN Open Source License
UMLS	www.nlm.nih.gov/research/umls/	open-source, free of charge, with license agreement. License agreements with other resources within UMLS must be arranged separately

6.2.1 Metrics

It is very difficult[8] to say something about the size of any of these terminologies, but we have made a table with the claims of their manufacturers.

[8] See the discussion under OpenGALEN for explanation why only stating numbers can be misleading

Terminology	Metrics
ID MACS	more than 500.000 different medical concepts
LinKBase	> 1 million concepts, >3 million terms in different languages
LOINC	32,000 observation terms
MED	over 59,000 concepts
MEDCIN	more than 215,000 clinical data elements
MESH	21,973 descriptors 132,123 Supplementary Concept Records 23,512 printed see references 102,346 other entry points
OpenGALEN	GRAIL is compositional, so this is not a trivial question to answer– there are at least three questions: • How many elementary entities (concept representations) are there? • How many sanctions have been added, implying how many composite entities which could be made? • How many entities have been made and named for one purpose or another by the current applications, mappings, etc.? We can give numbers for the number of asserted elementary entities and sanctions, but the more useful indication is the scope that is covered: at present, there is roughly half the coverage that we would eventually expect to find. However, there are important caveats: firstly, the depth and complexity in those areas that we have covered is greater than would be found in existing coding and classification schemes, and secondly, the system as a whole is well-developed and usable in specific areas. For example, the coverage of surgical procedures is extensive and virtually complete.
SNOMED CT	344.000 concepts organised into hierarchies more than 913.000 descriptions, 1.385.000 relationships
UMLS	875,255 concepts and 2.14 million concept names in over 100 biomedical source vocabularies, some in multiple languages

7 Results and Discussion

We here describe the main findings from our inventory. The main focus is on content. Though with the introduction of large detailed terminologies major investments in software are needed as well, this study did not look into specifications for and capabilities of software tools and delivery mechanisms for these terminologies. A follow-up study should address implementation issues in more depth. We did however make some first estimates on costs involved for software implementation.

7.1 Terminologies considered

7.1.1 The Case of SNOMED

The breadth and scope of SNOMED CT includes the content required for electronic medical records. The SNOMED CT Core terminology provides a common language that enables a consistent way of capturing, sharing and aggregating health data across specialties and sites of care. SNOMED International maintains the SNOMED CT technical design, the core content architecture, and the SNOMED CT Core content that includes the concepts table, the descriptions table, the relationships table (hierarchical and other semantic), a history table, mapping, and related technical documentation.

Originally in this study it was planned to look inside SNOMED CT extensively. Since the $5000 costs for an evaluation version of SNOMED appeared to be out of budget, a detailed evaluation of the product has not been performed. In addition time constraints would also have prohibited a thorough in depth evaluation. We fortunate to have the kind support of both CAP, and some SNOMED users to acquire some global insights.

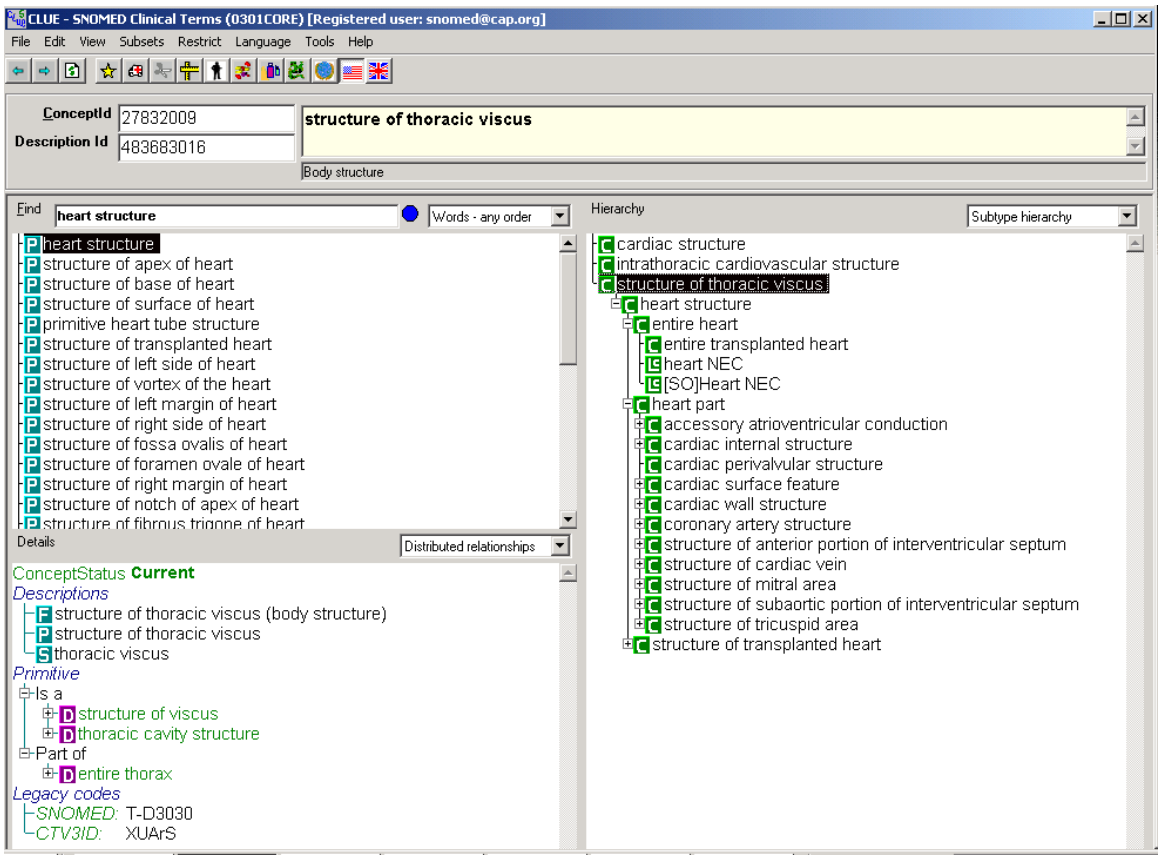

Figure 1. The CLUE browser showing some of the structure of SNOMED CT focussed around the concept Heart.

From interviews we understood that such an evaluation of SNOMED CT might even be at odds with the license conditions. Also the terms and conditions of the SNOMED CT license agreements are confidential[9]. So discussing conditions and terms of delivery here is also impossible on those grounds. It is even impossible according to the rules to jointly evaluate the product with licensed third parties[10]. In our discussions with the SNOMED organisation we learned that the licensing scheme and the conditions are under discussion. As a result terms and conditions are likely to change in the near future, to which extent is unknown at present.

During our interviews, the question rose 'What in essence is SNOMED CT?' According to the HIPAA NCVHS[11] Patient Medical Record Information Terminology questionnaire it is a set of ASCII tables delivered on CDROM or via download. In addition some example software algorithms are provided. The nature of that documentation falls under the SNOMED CT License Agreement and is hence unfit for public assessment[12]. Curiously in the HIPAA questionnaire it is explicitly stated that there is no need for any specific third party tools or products. Without appropriate delivery software the resource would be

[9] This situation has changed after our inventory, terms and conditions are public now

[10] This situation may have changed as well with the new licensing scheme.

[11] The National Committee on Vital and Health Statistics

[12] In a comment on this study CAP has stated it is public now.

rather useless, and not bring the advantage of being description logic based. Most other vendors do distribute their terminologies with software.

7.1.1.1 Formalism

SNOMED CT has been positioned as a third generation terminology, which means that it is based on a computer implementable formalism, in the case of SNOMED this is description logic. However, for the end user this is rather irrelevant, because SNOMED is delivered in the form of three text files, containing (1) concepts, (2) descriptions and (3) relationships. In principle, the contents of these files could be used to reconstruct the description logic representation, but unfortunately the specification of the description logic (and thus the classifier engine) is proprietary knowledge of a third party (Apelon). From personal communication with a SNOMED International representative early 2001, we concluded that even SNOMED International does not have access to the exact specification of the description logic under SNOMED. In the end, it seems that SNOMED CT as delivered to the end user is not that different from previous versions of SNOMED that weren't based on description logic.

7.1.1.2 Post-coordination

One of the features of SNOMED CT is called post-coordination. Contrary to a pre-coordinated classification, in which all possible concepts are pre-defined, in a post-coordinated classification new concepts may be composed from already existing concepts. The meaning of the term *post-coordination* in SNOMED is however limited to *composing*.

As an example (copied from the SNOMED slide http://www.snomed.org/Users_group/Presentations2002/David%20Markwell/img026.jpg), consider the following:

Appendectomy + Using = Laparoscope

which is equivalent to the existing pre-coordinated concept

Laparoscopic appendectomy

which is a subtype of

Laparoscopic procedure

Although, it is possible to compose the post-coordinated concept [Appendectomy + Using = Laparoscope], there is no means to deduce that this is the same as the pre-coordinated concept [Laparoscopic appendectomy], and thus you can't get to the supertype [Laparoscopic procedure] from the post-coordinated concept. Suppose, the concept [Laparoscopic procedure] is a trigger for a protocol, then the post-coordinated concept will **not** trigger that protocol.

Generally, a different definition is used for *post-coordination*, it is not just *composition* but also *classification*. This means that a *post-coordinated* concept is classified into the existing classification, and it also means that the post-coordinated concept [Appendectomy + Using = Laparoscope] would **really** be the same as the existing concept [Laparoscopic appendectomy], and would thus be a subtype of [Laparoscopic procedure]. In other cases, a new post-coordinated concept may create a completely new view on the classification, and thus *coordinate* the classification, without effecting the already existing views. This latter interpretation of post-coordination is for example implemented in OpenGALEN.

7.1.1.3 Content

As said we have not looked into SNOMED CT in great detail. We did however have a short inspection of the release of mid 2002. We did find quite a few improper subsumptions and inconsistent constructs. In the editorial process of SNOMED there is mention of a authoring style guide. Still there was considerable diversity in the kind of descriptions used for modelling concepts. (A similar diversity can be observed in the surgical procedure dissections in GALEN). Apparently a lot more work is required for gaining unity. SNOMED does release new versions every 6 months. We have no insight of how adequate the update of reported deficiencies is (and how adequately users report deficiencies!). Our contact informed us that most of what they had reported was confirmed as already fixed in the next release. The mechanism of correcting/updating needs to be part of a more in-depth product evaluation.

7.1.1.4 Messaging

Along with a large number of other classifications, SNOMED is registered as a coding system with HL7.

SNOMED contains mappings to a growing number of other coding systems and terminologies (e.g. CTV-3, CDT-2, HHCC, ICD-9-CM, ICD-10, ICD-O, ICPC, LOINC, NIC, NANDA, PNDS, OMAHA System, and OPCS-4).

7.1.1.5 EPR

SNOMED concept codes are arbitrary 32-bit numbers (9 positions in decimal representation). Using qualifiers and post-coordination may result in combinations of such codes, resulting in codes of 20 or more characters, with only a practical (but no theoretical) limit in length. Storage of such arbitrary-length codes must hence be facilitated. As the codes themselves are without meaning (there is no way of determining meaning from a code other than lookup), preferably the terms related to the codes are also stored in the EPR.

7.1.1.6 Training

Using SNOMED involves getting acquainted both with the clinical content, as well as with the applications involved to use SNOMED. Courses provided by SNOMED are:

- SNOMED CT® in Today's Practice of Pathology
  Upon completion of this course, participants should be able to:
  - Identify the purpose, structure and content of
  - SNOMED CT as it relates to pathology and cytology data
  - Utilize SNOMED CT in everyday practice
  - Communicate with pathology information system developers to clarify their needs for recording and retrieval of computerized data

- SNOMED® CT Implementation Training Class Outline
  By the end of the course the attendee will:
  - Have a working knowledge of SNOMED and its features
  - Be able to identify and implement methods for data collection, data exchange and data retrieval using SNOMED
  - Have the knowledge and techniques required to incorporate SNOMED into a health care application

Other implementation support tools and services available through SNOMED are:

- SNOMED CT Core Content with Browser and Technical Reference Guide
    - Updates to the Core (data tables in SNOMED CT structure delivered twice per year n CD-ROM format)
    - Language Services (reserved for countries requiring translation)
    - Extensions to the Core (such as U.S. Drug Extension)

- Technical Implementation Guide
    - 170 + page handbook which provides practical insight into SNOMED CT components such as concepts, descriptions and relationships, navigation concepts and links, importing and updating SNOMED CT, text searches, data retrieval, mapping services, managing legacy data, hierarchial navigation, extensions, and more.

- Developer Took Kit
    - Indexes and tables to streamline the development process

- Users Guide
    - Essential resource for the end user that describes the inner workings of SNOMED CT, its design and contents without all the technical detail.

- ANSI Structure Standard
    - Provides the complete detail of the ANSI (American National Standards Institute) standard.

- Member Name Space
    - Enables an organization to build works, such as those listed below that complement the SNOMED CT core to meet local needs while avoiding the potential for conflicting identifiers in HL7 messages.

7.1.2 Other systems

In principle at least some of the other systems could be considered as a candidate for a national clinical terminology in The Netherlands. We have not specifically addressed the suppliers with the question if their product is suitable for such a role and if they were willing to consider that, and what conditions would apply. Before adoption, a number of criteria must be satisfied (see 4.4 Deployment). Some of the criteria are stricter than others. Some systems have proven to be too academical to consider. Others do not have a clear open editorial process. With possibly the exception of SNOMED CT with its adoption by US and UK governments, none of the systems are near the situation of being considered an international (defacto) standard clinical (reference) terminology. We must admit at this point that there has been insufficient resource available for a sufficiently deep assessment of the other systems.

Apart from the other systems enumerated in the appendix, there is an alternative to consider. As mentioned under the ongoing work in Australia (see 6.1.1) one might develop mini-reference terminologies out of the existing classifications. To some extent that has also been done for the new French procedure classification CCAM. This

classification has been dissected in a specific categorial structure. A categorial structure can be loosely described as an ordered set of concepts, links, and rules that prescribe how these links and concepts may be combined. Each class in a classification is described independently in the form of a small semantic net representation. Describing existing classifications in a categorial structure is a good starting point for further formalisation in e.g. description logic based systems.

7.2 Economical aspects

It is difficult to estimate the costs of introducing a clinical terminology in the Netherlands, it is not just the terminology itself, but also training and software development. To give a rough idea the license costs for SNOMED CT are given, based on a preliminary quote we received from the SNOMED International management at the end of April 2003.

7.2.1 License costs for SNOMED CT

The SNOMED International Authority refined its strategy with national governments in April as planned. It will now make its proposal to the CAP Board; therefore any proposal will be subject to the Board's final review and acceptance.

The SNOMED National Government Charter license fees are tiered based on the size of the population in the country. This level is designed for use at local, regional and national applications within a country. License covers all government services, all government employees wherever located, all contractors for government project purposes and anyone reporting to government within the country. The fee would enable use of the SNOMED CT Core terminology and infrastructure, enable covered entities to create SNOMED compliant works, and to participate in the National Governments Advisory Group as well as numerous working groups related to content, technology and implementation. The upfront fee supports some of the start-up activities such as translation/localization, mapping starter sets, training of the country's core team at CAP headquarters, national name space to accommodate extensions and the availability of a dedicated support person.

With the Netherlands population of about 16M, the National Government Charter license fee would be comprised of two parts:

- A base royalty of $2.3M[13];
- An annual use royalty of $850,000.

Commitments for a minimum of 5 years would be required. Longer term arrangements are possible. Fees escalate 4% each year, this means that the annual use royalty will be about $995.000 at year 5, and the total license costs over the first 5 years $6.9M.

7.2.1.1 Additional costs

Included in the SNOMED CT license are the translation costs of an approximately 20-25 person/year effort. Apart from the license costs for the SNOMED sources there are additional costs involved to implement it in the clinical environment. Terminology

[13] This royalty includes translation. In their comments CAP stated a base royalty of $1.4 for the original US/UK language version.

services must be developed and/or purchased. Clinical systems must be adapted. For Clinergy, a GP software product based on an early version of GALEN, it appeared that around 1996 50% of their effort was in terminology development, and 50% in making the terminology accessible via software. It is difficult to say if these figures scale to multiple systems adapting to one single terminology. In our interviews we did not get any hard figures. Nevertheless we think that if the Dutch software industry is going to implement in the next 5 years, some 5-25 Million Euro will be spent, depending on the level of sophistication.

Extrapolating UK figures to the Dutch situation one should consider additionally a 1-2 Million Euro annual expenditure for maintenance and training.

Additional costs presented here under the SNOMED heading are likely to be valid even when another system is chosen.

7.2.2 New developments

Looking at the costs of purchasing a license for SNOMED CT, it is still worth considering the possibility of developing a new national terminology for the Netherlands. The costs of the annual use royalty of $850.000 will support an organization with at least 6 full-time employees. It also seems reasonable to assume that the annual costs will decrease once the terminology reaches a mature state. It can also be questioned if the size of SNOMED CT is really needed, or that one is better off to gradually grow a terminology on the basis of end user needs. Such an approach has been taken in the development of the Read 3 system. That development was seriously hampered by the lack of proper supporting technology at that time. One could seriously reconsider that on the basis of new insights.

A serious disadvantage might be that a local development leaves the Netherlands outside a possible world standard. It also restricts the market for the export of Dutch software applications to other markets. Mapping a new terminology to existing classifications, like the ICD family, or including it into the UMLS may reduce this risk. The UMLS also includes mappings to SNOMED. A definite advantage is growing national expertise in advanced clinical terminology development. It also brings new business chalenges to Dutch and European industries. Also the intellectual property rights on the terminology stays with the Netherlands.

Extending the Dutch-only option would be the development of a new terminology for the use within the European Union. The total costs for annual royalties of SNOMED for every national government within Europe could reach $20M per year and including the base royalty $160M over the first five years (based on 370 million people in the current European Union).

7.3 Introduction strategies

There are several strategies possible for the introduction of a national clinical terminology.

An obvious solution for rolling out a clinical terminology (for instance SNOMED CT) is to let the marketplace decide. Software vendors could obtain sub-licensing rights from SNOMED International and collect license fees from their customers. This model will

stimulate that the product is only bought if it is really useful and useable. In the Dutch marketplace some 400 companies are active in health ICT. Only a limited number are involved in development and sales of terminology related products.

Both the vendors of terminology as well as the vendors of terminology servers do have a keen interest in settlements for national licences. It greatly reduces their costs, and savings on a national scale can be up to 90% if compared with individual site licences. A national government license reduces the uncertainty and signals to vendors and end users that their investments for systems using the terminology would be more secure.

The second scenario is that purchase of a national licence, either through government funding or some other central resource. The product is made freely available to all Dutch interested parties, and one just sees what happens in product development. This seems more or less the scenario of the purchase of a national license in the USA by the National Library of Medicine. To stimulate adoption, some accompanying measures in the form of regulation, or agreements between parties in the healthcare sector should be in place.

A third scenario is to consider a few well-defined representative experiments to prove in real practice what is needed. For these experiments only a few site licences are needed. This scenario will result in more understanding of what is really needed, and what the consequences are for the care delivery process.

7.4 Conclusions

As outlined in Chapter 2:Objectives, the starting point of this study was the existence of candidate terminologies (most notably SNOMED CT), rather than starting with an in depth analysis of terminology needs.

Complex description logic based terminologies like SNOMED CT are still in their infancy. Behaviour of such complex systems is only partially understood. Debate on issues of terminology and concept representation is ongoing. There is no grand unifying theory for medical language, and hence we will have to live with the fact that we cannot represent all medical expressions in a coherent way. Yet we must also consider where we did make advances. Some systems have progressed to a stage where large-scale implementation seems feasible. Hence it is time for serious real life experiments before a properly motivated decision for a centrally financed licence for SNOMED CT can be made.

From our discussions with government and industry representatives it appears that they see quality control as the most immediate benefit of the introduction of SNOMED CT. Simple reminder systems based on Arden syntax[14] and similar are envisaged. In that setting a SNOMED CT based terminology server functions as a mediator between data present in the patient record, and facts represented in the rule(s) of the knowledge base. This scenario does require data entry at the point of care!

Users of terminology traditionally have not had much influence on the standard terminologies like ICD. They did however have major influence on their professional terminologies (often derived from official WHO classifications). With the emergence of

[14] Arden syntax is a rule-based knowledge representation paradigm. Its intended for one, or in some cases a few, rules. The syntax does support localisation of terminology in its rules.

international, (commercial) terminologies, the influence of individuals, and even national professional organisations will diminish, unless they are given a role in the production of subsets. Experience with derivatives of ICD has learned that not only subsets, but in addition (local) refinements of terms often are needed.

Whatever is decided about a clinical terminology for The Netherlands, it is obvious that substantial continued investment in medical terminology is needed. The complexity of systems and demand for frequent updates and multi-version control requires a well-staffed expertise centre. If the sector is properly organised substantial cost savings on national expenditure can be achieved. Policy makers not only at the national level, but also at the institutional level should start thinking on allocating resources.

Many experts in the field of clinical terminology have expressed that on financial grounds there would only be room for one resource for 'all of medicine'. We have shown that if the European Union chooses to develop its own open source healthcare terminology instead of individual countries purchasing national licenses for SNOMED RT, there would be ample funding for such an enterprise. Even though Europe has a better understanding of multi-linguality, is used to multi-cultural settings, and has given birth to technologically much more advanced terminology systems, it is highly unlikely that a joint European Health Language effort will ever emerge. Where terminology now is widely considered as a major infrastructural hurdle to take in the development of the next generation of health records, the EU, as most national governments, is absent in taking responsibility.

A short investigation on the mid 2002 release of SNOMED CT has left us with concern about present quality of the resource. Given the UK-USA dominance on the editorial board there is also serious concern about influence of other countries on the product. In further studies it must be investigated how well and how efficient (inter)national specialists groups can communicate with SNOMED International on their requirements. There also is serious concern about the great number of terms in SNOMED CT and the precision of descriptions/definitions of terms. This might lead to an increase in misclassification. We did not find inter-rater variability studies. Neither did we find those for the other systems considered. This will need further investigation.

It remains confusing what 'SNOMED out of the box' will do for you. SNOMED is presented in the context of description logics, but actually does not use or require a description logic driven software for its deployment. Another notion is that it is a reference terminology. This not only means that existing classifications like disease and procedure classifications are mapped to it, but also the texts seen in user interfaces of software products (=interface terminology) could be mapped by the software vendor to SNOMED CT. In practice this may lead to the situation that the user never sees a bit of SNOMED. The real essence of SNOMED is from the software engineering point of view a piece of middleware. It may support the mediation between health care actors, be it patients, carers, or software (knowledge agents). If it only were a reference, there seems no need for translation. If it is indeed both reference and interface terminology, than it remains to be seen what should be translated.

There is a danger to become too dependent of one or two suppliers. If there is the need to migrate to an other solution, access to historical data may be in danger. Some suppliers may demand all codes to be removed from systems at the expiration of the

license agreement. Proper agreements or additional provisions for continued data access are a must. As more countries are presently evaluating SNOMED, the suggestion has been made to set up an Independent International Cooperation for the Evaluation of SNOMED CT. The requirements worked out in paragraph 4.4 strongly suggest striving for some sort of open source licence agreement. A group of countries considering national implementation is in a much stronger position to negotiate with SNOMED International on that issue.

A serious alternative to consider is the national development of a clinical terminology starting from the existing WHO FIC classifications and the national procedure classifications. If such an activity is started, terms should at least be modelled in a well-defined categorical structure, which will facilitate conversion to more formal ontologies. Such work would also be useful as a validation of the existing SNOMED mappings that were made on the basis of traditional consensus methods.

8 Recommendations

The report addresses many issues. We have been critical on systems. We have concluded that there still is both at the supplier's side as well as at the users side a lack of sufficient understanding on the introduction of large terminologies in healthcare. Rather than recommend the adoption of a certain solution our recommendations are about how to move ahead to gain sufficient understanding, and the conditions that must be satisfied to make progress. This study is restricted to what the options are. The recommendations in this document are made to the Dutch healthcare sector as a whole. Who should do what and who should pay is outside the scope of this study. The following is presented for consideration:

The main objective of this study is to answer the question 'Is it sensible to aim for a national introduction of SNOMED CT, and if so for which sectors of healthcare?' The conclusion on the product SNOMED CT is that it has not yet reached sufficient maturity. At the same time we conclude that there is a lack of understanding on precise requirements of the sector. An immediate full-scale introduction will not guarantee useful and useable products within one or two years time. We expect this will even be counter-productive. We do recommend to start with a few learning experiments. At this time we do not have sufficient insight in the closed, proprietary systems to see how well they would fit into this. Most of them seem to act as an interface terminology to the existing reporting terminologies, while others (MEDCIN) seem to also cover all the detail of the clinical narrative. We suggest using SNOMED CT for those experiments, there are many publications about it, it is considered as de-facto standard, and it is large. Experiments should be done properly to allow for some generalisation to other options like OpenGALEN, LOINC, new developments, etc. We have noted two major classes of claims where detailed common terminology is said to give added value. These are in a) quality control and the management of healthcare delivery, and in b) communication between dissimilar healthcare organisations. The experiments should clarify the needs of the sector, and prove the added value of a formalized terminology like SNOMED CT.

Recommendation 1: *Do conduct at least two major real life learning experiments with SNOMED CT. Make on the basis of the outcome a decision on SNOMED CT as a national standard.*

The existing (inter)national classifications for disease, procedures, and others will be around for at least the next decades. There is general concern about the quality of data coded using these classifications. Such coded data is important for management (costing), health statistics (policy making), and epidemiology, to name just a few. In this era of evidence-based medicine we can not only build on the costly randomised control trials, but we also need quality in routinely collected observational data to extend medical knowledge. The reporting classifications will play a continued role in recording even after introduction of SNOMED CT. Further formal analysis of the classifications on the basis of sound ontological principles is necessary. This will ensure better mappings to formal systems. In some countries this analysis is seen as a possible starting point for gradually development of a clinical terminology.

Recommendation 2*: Continue with a more principled further development and maintenance of existing reporting classifications. Do investigate if this new approach is a viable start for a clinical terminology.*

Present practice on specialist terminologies in Dutch health care is that each professional body acts more or less on its own. Only occasionally a terminologist is involved in the construction of a set of terms for a certain purpose. There is neither a central reference resource nor central expertise to align the individual activities and products. The ultimate role of a clinical reference terminology is to enable detailed documentation of all patient cases, and to automatically derive data in reporting terminologies (ICD, CMSV (procedure), DBC (grouper), etc.). The traditional coding schemes are updated in cycles of years or even decades. A clinical terminology is updated typically twice a year. To maintain this central resource with its links to derivatives is a complex task for skilled people. From 1974-1995 the Dutch Standing Committee on Classification and Coding (WCC) has functioned as an expertise centre for all issues relating to medical language. At the end of the 80's beginning 90's several countries have adopted this model. Because of strong interdependencies of products, and increasing levels of complexity of terminologies, and the need for short communication lines, it is advised to follow the Australian model and concentrate the work on all terminologies (including procedures, diseases, and groupers) in one physical unit.

***Recommendation 3**: Establish a national expertise centre for all Health Terminology*

Recorded data should as much as possible be represented in a standardised form. To preserve space, linguistic expressions are often recorded in coded form only. At all times access to such data must be secured. Some proprietary systems do not allow continued usage after expiration of the license to use the product. Not even on historical data. Be aware that if a system is abandoned, also the relationships between individual terminological expressions are lost! No external circumstances should force users to abandon their terminology. There should be no constraints from the supplier to optimise the terminology for national use. Also at the sub-national level there should be adequate facilities to (temporarily) define term sets for e.g. intervention studies. These requirements demand sufficient editorial access. Terminology is an important element of the infrastructure that facilitates communication between all parties in the sector. It should therefore also be equally accessible to all. Appropriate access to the infrastructure seems best guaranteed under Open Source conditions.

***Recommendation 4:** Secure appropriate access to coded data. Insist on Open Source conditions for a national reference terminology*

Quality of narrative data depends on quality of the terminology used. The new and large terminologies are delivered to the clinical user via Terminology Servers. In some sense terminology has become software. Because of the sheer size and compositional nature of the new terminologies, it is impossible to enumerate and validate all possible expressions. There is a tendency to increasingly use terminological reasoning in systems that generate automatic reminders. Those systems do get the nature of a device, and may be considered so in the context of being accepted for clinical use. As it is very unlikely that there will ever be a system that does pass a formal proof of correctness, some other quality indicators against which a system can be certified are needed.

***Recommendation 5:** Further develop criteria for specification and certification of terminologies and derived products*

A good clinical terminology is central to the future clinical systems. Where terminology traditionally has been the work of many volunteers, it is now gradually turning into a mature industry. This means that quality products should be expected, which do cost real money. We have indicated for the next five years that the annual costs for just a national license is over 1 Million Euro. The adoption, implementation, and maintenance of a broad clinical terminology bears substantial cost in comparison to what is spent presently on terminologies. It is hard to get precise figures, because our major reference is on some facts about, and assumptions made on other healthcare systems (US-UK-AU). Nevertheless on present knowledge it is safe to say that for a country like The Netherlands, whatever approach is chosen a basic recurrent annual expenditure of 1-2 Million Euro is expected. Next to these basic central costs, there is an additional cost for implementation by industry, which on first estimate may range from 5-25 Million Euro in the next 5 years. There is insufficient awareness about these potential costs. It is as yet unclear how the central investments could be covered.

Recommendation 6: *Be prepared for a central annual expenditure of at least 2 Million Euro on health terminology*

9 Appendix A: Justification and Acknowledgements

The conclusions and recommendations in this report were based on a short inventory project conducted at the Department of Medical Informatics, University Medical Centre Nijmegen, with support of the Department of Medical Informatics, Amsterdam Medical Centre. The authors from Nijmegen were core members of the GALEN programme, and played a central role in GALEN software development. OpenGALEN is one of the products considered in this study.

We have spoken to a number of stakeholders around the world on different issues. Their names and roles are listed below. We appreciate that they have shared their time and knowledge. The inclusion of their names is in no way meant to be an endorsement of the conclusions of this report. The conclusions and recommendations are those from the authors.

Country	Organisation	Contacts	Issues
AU	Dept Health and Aged	Peter McIsaac	Deployment in primary/secondary care situation
AU	OpenEHR	Sam Heard	Role of terminology in advanced record systems
NL	AMC	Henk Lamberts	Relationship with ICPC-ICD
NL	AMC	Nicolette de Keizer	DICE
NL	AMC	Kitty Jager	European Renal Registry
NL	AMC	Anita Ravelli	Perinatal Care
NL	AZG	G. Stallinga	Continuity of Care (CVA-project)
NL	HL7	Bert Kabbes	
NL	MediLingua	Simon Andriesen	Translation methodology, economics
NL	Npi	Dr Y.Heerkens	Relation with ICF (formerly ICIDH)
NL	Prismant	Wim Wilhelm	Relation with existing SNOMED2 based national pathology registry
NL	Torex-Hiscom	Jacob Hofdijk	Implementation in Legacies, links with costing (DBC)
NL	UMCU	Loek Leenen	Data-intensive registration in traumatology
NO	KITH	Glen Thorssen	Norwegian planning of SNOMED introduction
SE	SIS	Gunnar Klein/ Goran Holmberg	Exchange of experiences en results of Dutch and Swedish SNOMED evaluation projects

UK	CHIME	Dipak Kalra	Role of terminology in advanced record systems
UK	G7-project on medical terminology	Ray Rogers	Worldwide stakeholders
UK	Isoft	W.D. Solomon	SNOMED implementation issues for HIS suppliers
UK	Lloyd-Nolder Associates	Colin Nolder	Change management
UK	NHS-Information Authority, Birmingham	Colin Price, Edward Cheetam	Deployment, maintenance, version control, quality control
UK	OpenGALEN	Alan Rector Jeremy Rogers	Development, maintenance, formal properties
UK	Salford University	Nick Hardicker	Nursing terminology
UK	Royal College of Nursing/SNOMED Editorial board	Anne Casey	Nursing terminology
UK	Centre for Health Services Research /SNOMED Editorial board	Nick Booth	Editorial issues, version management, user involvement
UK	SCHIN	Neil Jones, Peter Johson, Richard Hall	Prodigy project; linkage to knowledge
USA	Kaiser Permanente	John Mattison	Clinical situations which demonstrate added value of SNOMED deployment
USA	SNOMED/CAP	Diane Aschmann	Usage of SNOMED product, product information

10 Appendix B: Bibliography

The documents below have been used as background for this report. Most of the documents are internal documents, or documents downloaded from WEBsites.

1. SNOMED® RT Backgrounder, CAP, 07-18-2000
2. About SNOMED® International, CAP, 07-18-2000
3. Paper Health Records vs. the SNOMED®-Coded Electronic Health Record, CAP 07-18-2000
4. Quality Assurance at SNOMED® International, CAP 07-18-2000
5. SNOMED International: Our Commitment to Health Care, September 2002 Meeting
6. SNOMED® INTERNATIONAL INTRODUCES THE MOST COMPREHENSIVE
7. HEALTH CARE TERMINOLOGY AVAILABLE WORLDWIDE, Press Release CAP, February 13, 2002
8. SNOMED® Use Agreements
9. SNOMED® Clinical Terms™ Technical Specification, Word Search Tables, Version 6, 2001-09-24
10. SNOMED® Clinical Terms™ Technical Specification, Canonical Table, Version 2, 2001-09-25
11. SNOMED® Clinical Terms™ Technical Specification, Roadmap to SNOMED CT Documents, Version 2, 2000-10-12
12. SNOMED® Clinical Terms™ Technical Specification, Core Structure, Version 23, 2001-12-12
13. SNOMED® Clinical Terms™ Technical Specification, Relationships to Earlier Terminologies, Version 1, 2001-12-12
14. SNOMED® Clinical Terms™ Technical Specification, Glossary, Version 3, 2001-12-12
15. SNOMED® Clinical Terms™ Consultation Document, Proposed Structures for
16. Cross Mapping Tables, Version 5, 2000-10-12
17. SNOMED® Clinical Terms™ Consultation Document, Requirements Analysis, Version 10, 2000-10-12
18. SNOMED® Clinical Terms™ Consultation Document, Subset Examples, Version 8, 2000-10-12
19. SNOMED® Clinical Terms™ Consultation Document, History Tables and Mechanisms, Version 6, 2001-12-12
20. SNOMED® Clinical Terms™ Consultation Document, Subset Tables and Mechanisms, Version 19, 2001-12-12
21. SNOMED® Clinical Terms™, Requirements Document, Proposals for the New Terminology, Version 5, 2000-06-03
22. SNOMED® Clinical Terms™, Discussion notes, Implementation Guide Outline, Version 1, 2000-10-12
23. SNOMED® Clinical Terms. Content Specification, Concept Hierarchy: Upper Levels, Version 004, 2001-07-03
24. SNOMED® Clinical Terms. Concept Hierarchy: Upper Levels – Questionnaire, Version 1, 2001-07-20

25. SNOMED® Clinical Terms™, Consultation - Subset Mechanisms, Questionnaire, Version 0.1g(US) July 2000
26. SNOMED Clinical Terms® (SNOMED CT®), January 2003 Release, Fact Sheet
27. NHS Acceptance Criteria for SNOMED® Clinical Terms™, NHS, May 2002
28. A World of Possibility, SNOMED International, SNOMED CT Brochure, September 2002
29. SNOMED Clinical Terms, An update on progress, Paul Iggulden, Powerpoint presentation at HC, 19th March 2002
30. Representing Clinical Data in the Healthcare Enterprise, Use of SNOMED® in Integrating Clinical Viewpoints, Michael Stearns, Powerpoint presentation to CEN TC251, 4 February 2001
31. SNOMED® Clinical Terms, Technical Design, David Markwell, Powerpoint presentation to CEN TC251, 4 February 2001
32. SNOMED RT: A Reference Terminology for Health Care, Kent A. Spackman, Keith E. Campbell, Roger A. Côté, AMIA Fall Symposium Proceedings, 1997
33. Building and Validating SNOMED®RT: A Preview of First Release, Kent A. Spackman, November 5, 1999
34. Successful Search and Retrieval using SNOMED Clinical Terms, Kent Spackman, David Markwell, http://www.snomed.org/Users_group/Presentations2002/David%20Markwell/ , 8 November 2002
35. Clinical Terminology: Why is it so hard?, Alan Rector, Methods of Information in Medicine, 1999
36. Supporting the use of the GALEN Intermediate Representation, W.D. Solomon, J.E.Rogers, A.L. Rector, E.J. van der Haring, P.E. Zanstra, Annual Fall Symposium of the American Medical Informatics Association, 1998
37. GALEN: a third generation terminology tool to support a multipurpose national coding system for surgical procedures, B. Trombert-Paviot, J.M. Rodrigues, J.E. Rogers, R. Baud, E. van der Haring, A.M. Rassinoux, V. Abrial, L. Clavel, H. Idir, International Journal of Medical Informatics, 2000
38. Health informatics - Electronic healthcare record communication, partI-IV, ENV-13606, CEN
39. Health informatics - Categorial structure for classifications and coding systems of surgical procedures, EN-1828, CEN
40. Archetypes Constraint-based Domain Models for Futureproof Information Systems, Thomas Beale, www.deepthought.com.au, 2001
41. Background and overview of the Good Electronic Health Record (GEHR), Dr Peter Schloeffel, Thomas Beale, Dr Sam Heard, Dr David Rowed, openEHR Foundation, May 2001
42. Validating Electronic Health Records Using Archetypes and XML, Tun Z., Bird L.J. and Goodchild A., titanium.dstc.edu.au., 2002.
43. GALEN 10 Years On: tasks and supporting tools, Jeremy Rogers, Angus Roberts, Danny Solomon, Egbert van der Haring Christopher Wroe, Pieter Zanstra, Alan Rector,
44. Rubrics to Dissections to GRAIL to Classifications, JE Rogers, WD Solomon, AL Rector, P Pole, P Zanstra, E van der Haring
45. Language, Intermediate Representations, and Expressivity: Issues in developing an ontology for medicine, AL Rector, Jeremy Rogers, Danny Solomon, Werner Ceusters, Pieter E. Zanstra, Egbert J. van der Haring, Angelo Rossi Mori

46. A comprehensive approach to developing and integrating multilingual classifications: GALEN_s Classification Workbench, Alan L. Rector,M.D. Ph.D., Werner Ceusters, Jeremy Rogers, Angelo Rossi Mori, Egbert van der Haring, Pieter Zanstra
47. Using the GRAIL language for Classification Management, P.E. Zanstra, F. Flier, E.J. van der Haring, J.E..Rogers , W.D. Solomon
48. Bridging the gap between developers of terminologies and record architectures, Pieter E. Zanstra; Final report of the HCT 3108 TOMELO project, 1999 http://www.mi-kun.nl/tomelo
49. Consultancy to develop the requirement for achieving an Australian general practice terminology suitable for recording at the point of care, FINAL REPORT, General Practice Branch Commonwealth Department of Health and Aged Care, Elizabeth Moss, November 2001
50. Options for Australian health terminology development, NCCH, April 2002
51. Policy Issues in Health Terminology Standards for Computerised Records, Jeremy Firth, June 2002
52. AUSTRALIAN POLICY ON HEALTH REFERENCE TERMINOLOGY, Recommendations to National Health Information Standards Advisory Committee from Terminology Group, General Practice Computing Group, September 2002
53. An approach to terminology standards for electronic health records: conceptual framework and issues, Siaw-Teng Liaw Heather Grain, February 2002
54. SNOMED CT - Anaesthetic and EPR Perspectives, Dr Roger Tackley, Powerpoint presentation, Januar 2003
55. SNOMED CT ALPHA TRIALS: Report from Australia, Rosemary Roberts,Peter MacIsaac, Powerpoint presentation, April 2003
56. Campbell, K. E., S. P. Cohn, et al. (1998). "Scalable methodologies for distributed development of logic-based convergent medical terminology." Methods of Information in Medicine 37(4-5): 426-39.
57. Chute, C. G., P. L. Elkin, et al. (1999). "Desiderata for a clinical terminology server." Proceedings / AMIA Annual Symposium: 42-6.
58. Cimino, J. J. (1998). "Desiderata for controlled medical vocabularies in the twenty-first century." Methods of Information in Medicine 37(4-5): 394-403.
59. de Keizer, N. F., A. Abu-Hanna, et al. (2000). "Understanding terminological systems. I: Terminology and typology." Methods of Information in Medicine 39(1): 16-21.
60. de Keizer, N. F. and A. Abu-Hanna (2000). "Understanding terminological systems. II: Experience with conceptual and formal representation of structure." Methods of Information in Medicine 39(1): 22-9.
61. ISO/TC215, W. (2000). Standard Specification for Quality Indicators for Controlled Health Vocabularies, ASTM Committee E31.
62. OMG (2000). Lexicon Query Service Specification, Object Management Group.
63. King, M. S., Lipsky, M. S., and Sharp, L. Expert agreement in Current Procedural Terminology evaluation and management coding. Arch Intern Med 162(3), 316-320. 11-2-2002. 14-3-2002.
64. Lu, T. H., T. P. Shih, et al. (2001). "Diversity in death certification: a case vignette approach." J Clin Epidemiol 54(11): 1086-93.

11 Appendix C: Other systems

11.1 ID MACS

ID MACS 1.0 – medical semantic network, the result of merging the works of Friedrich Wingert and ID DIACOS, is a revolutionary product for the use of medical terminology in coding systems and other IT health case solutions.

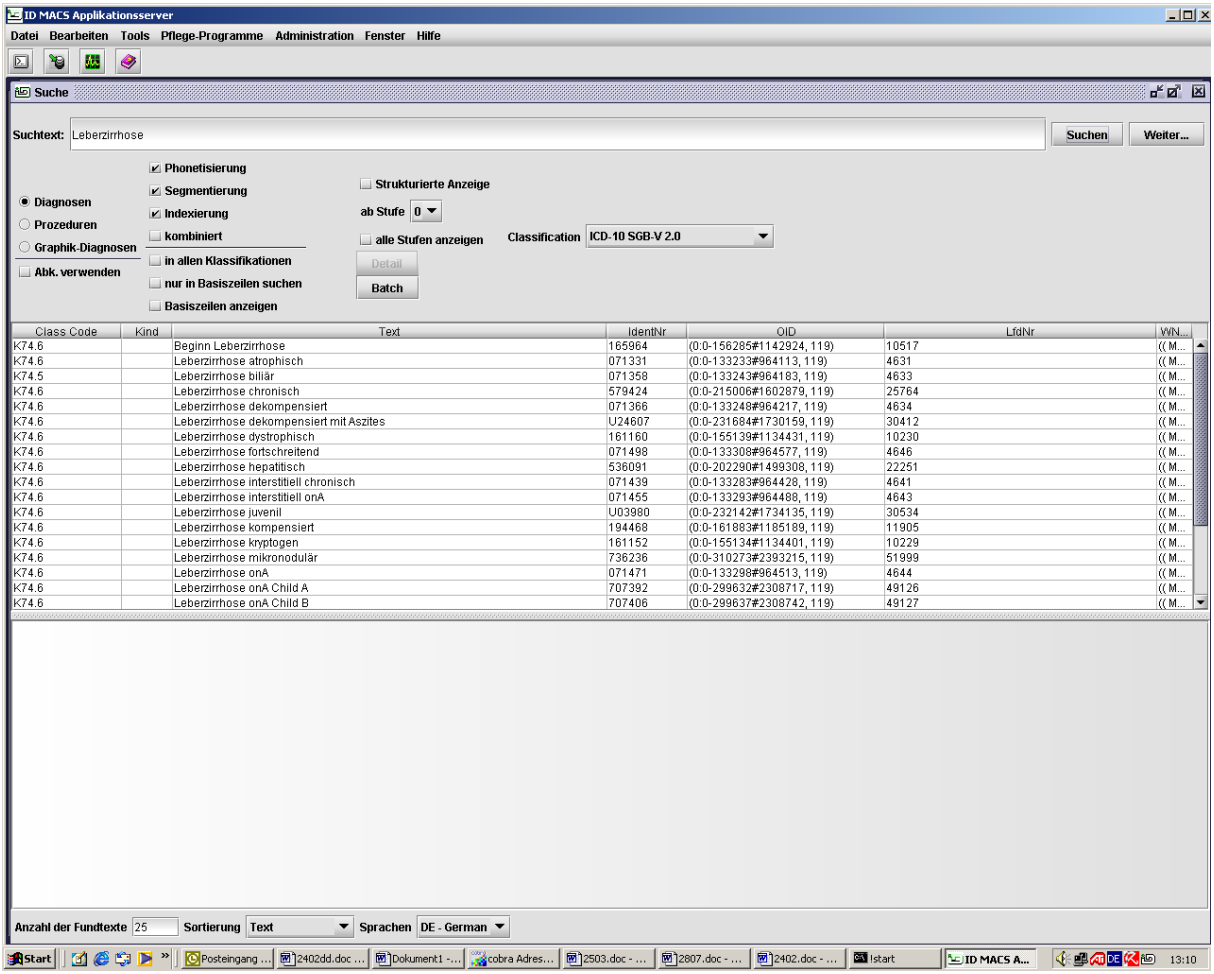

The entire terminology of medicine and nursing has been defined according to segmentation and rules of syntax and the terms have been arranged within an axis model of ten dimensions structured multilingually. Due to its interaxial and intraaxial relations, ID MACS can be implemented in many fields like process management, coding of diagnoses and procedures, and reimbursement to name but a few.

The axis model's view can be pre-set with regard to the specific requirements of clinical studies, coding guidelines and billing. ID MACS is working via an object-oriented database ensuring quick access to other knowledge bases like classifications or medical literature.

Multilingually designed ID MACS – medical semantic network provides links to classifications, thesaurus models and nomenclature, which helps to catalyse the development of IT solutions in medicine, nursing, epidemiology and public health. It enables an efficient way to access data for health reports, clinical studies as well as reference databases.

11.2 LinkBase

LinKBase® is an expansive medical knowledge base that provides the "intelligence" used by all L&C applications. In fact, it is the world's largest formal medical ontology, i.e. a conceptual computer-understandable representation of medicine. Due to its magnitude, formal structure and the fact that it is machine-readable, LinKBase® is the only medical knowledge base capable of producing the results needed in automated processes that work with medical unstructured texts.

LinKBase® is unique in the marketplace as it has been continuously developed over 14 years creating a depth in the semantic linkages between concepts and a state of fine-tuning that can only come with time and dedicated experts. Another distinguishing feature is that LinKBase® was built from the beginning with NLP in mind in contrast to other medical ontologies currently available.

Medical language is used in a variety of different ways. Medical professionals in different locations, or even within the same hospital, use different expressions to indicate the same medical idea. So, if human beings don't completely understand each other, then maybe computers could be designed to understand medical language and process it in the most efficient and meaningful way. Just supporting medical data processing, classification systems as ICD-9, SNOMED, et al will not suffice.

Current coding classification systems have thousands of terms (e.g. about 20,000 for ICD-9-CM) in a one-dimensional hierarchical structure. However, for a computer to understand and process medical expressions, it needs a complete, multi-dimensional representation of the medical world where concepts are linked to each other via several types of relationships. This representation is called an ontology.

Language and Computing has developed a technology that breaks down medical language to a common denominator set of medical concepts that can be expressed both by standardized terminologies and by natural language expressions.

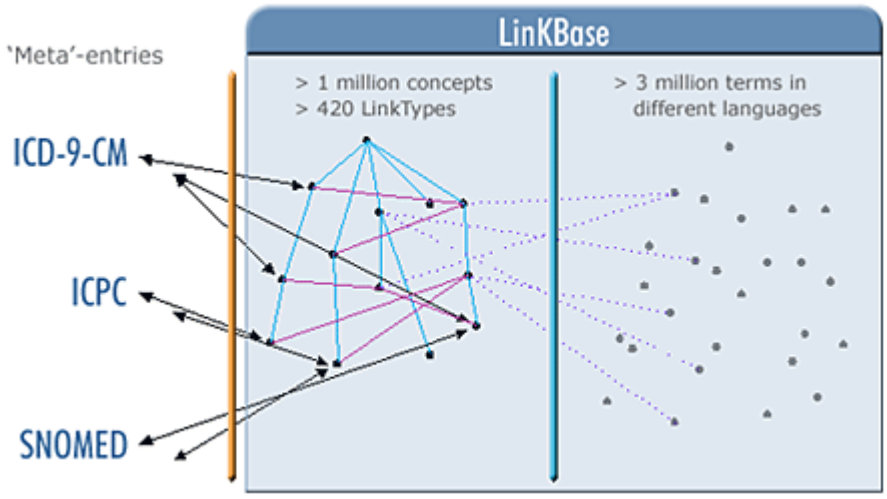

LinKBase® Ontology is a formal conceptual description of the medical domain and as such is rendered machine-readable by a computer.

The medical concepts are language independent and are linked to about 3,000,000 terms in various languages (English, Spanish, French, etc.). Terms can be stored in different languages and can be linked to concepts, criteria and link types with an intersection table allowing L&C to define both homonyms (1 term that has several different meanings or linked concepts / criteria / link types) and synonyms (multiple terms associated with 1 concept/criteria/link type).

LinKBase® is the powerful foundation that drives all other L&C applications. It enables cross-mappings from one classification system to another (i.e. ICD-9-CM, ICD-10, SNOMED, Read, MedDRA, etc.); it accommodates searches in one language and information retrieval in another language; it optimizes indexing of medical documents based on their content, and it facilitates the extraction of medical facts, etc.

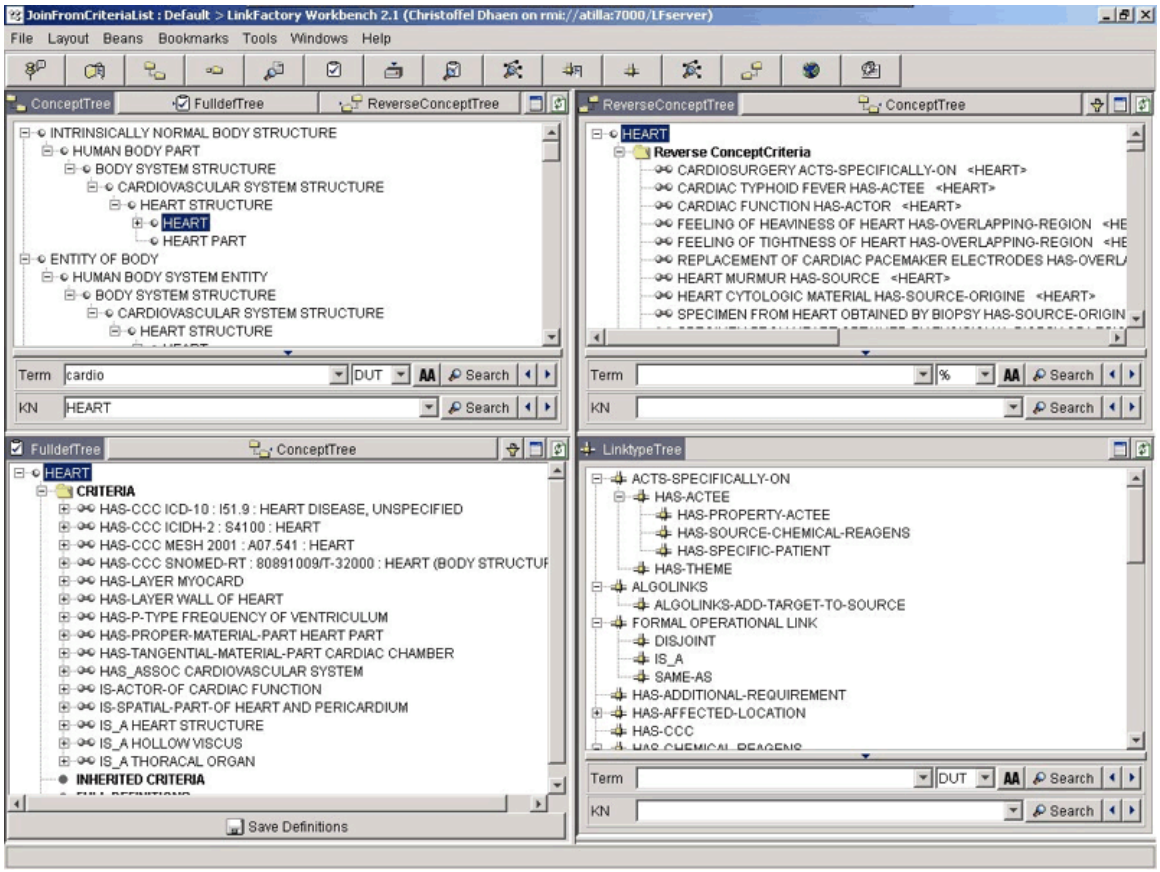

The LinKFactory® tool provides an effective and user-friendly way to create, maintain and extend extensive multilingual terminology systems and ontologies (English, Spanish, French, etc.). It can even be applied outside the medical domain and has been successfully tested. Since building a comprehensive and solid terminology system is very time-consuming and labor-intensive, it is essential to have the right tools to reduce the workload of this kind of project. LinKFactory® is very powerful, yet easy to learn and manage.

11.3 Logical Observation Identifiers Names and Codes (LOINC®)

11.3.1 Overview

The purpose of the LOINC database is to facilitate the exchange and pooling of results, such as blood hemoglobin, serum potassium, or vital signs, for clinical care, outcomes management, and research. Currently, most laboratories and other diagnostic services use HL7 to send their results electronically from their reporting systems to their care systems. However, most laboratories and other diagnostic care services identify tests in these messages by means of their internal and idiosyncratic code values. Thus, the care system cannot fully "understand" and properly file the results they receive unless they either adopt the producer's laboratory codes (which is impossible if they receive results from multiple sources), or invest in the work to map each result producer's code system to their internal code system. LOINC codes are universal identifiers for laboratory and other clinical observations that solve this problem.

The laboratory portion of the LOINC database contains the usual categories of chemistry, hematology, serology, microbiology (including parasitology and virology), and toxicology; as well as categories for drugs and the cell counts you would find reported on a complete blood count or a cerebrospinal fluid cell count. Antibiotic susceptibilities are a separate category. The clinical portion of the LOINC database includes entries for vital signs, hemodynamics, intake/output, EKG, obstetric ultrasound, cardiac echo, urologic imaging, gastroendoscopic procedures, pulmonary ventilator management, selected survey instruments, and other clinical observations.

The Regenstrief Institute (www.regenstrief.org) maintains the LOINC database and its supporting documentation.

11.3.2 Background

LOINC® is a voluntary effort housed in the Regenstrief Institute for Health Care, an internationally respected non-profit medical research organization associated with Indiana University. LOINC system was initiated in 1994 by the Regenstrief Institute and developed by Regenstrief and the LOINC committee as a response to the demand for electronic movement of clinical data from laboratories that produce the data to hospitals, physician's offices, and payers who use the data for clinical care and management purposes.

- The LOINC® laboratory terms set provides a standard set of universal names and codes for identifying individual laboratory and clinical results.

- LOINC codes allow users to merge clinical results from many sources into one database for patient care, clinical research, or management.

- The LOINC database currently contains about 32,000 observation terms.
- Nearly 20,000 of these observational terms relate to laboratory testing.

- Each record in the LOINC database identifies a clinical observation and contains a formal 6-part name, a unique name for tests identifying code with check digit, synonyms, and other useful information.

- LOINC records apply to all tests with equivalent clinical results. They are not unique per company.

- Distinct LOINC codes are required for each specimen for which your test kit has been calibrated. If your instrument/kit produced one value for each specimen and you recommend its use on two specimens-- say whole blood and CSF-- two LOINC codes are needed, one for whole blood and one for CSF. If two or more results per specimen are reported (e.g., a control value or a total and a percent), two or more LOINC codes are needed per supported specimen.

- LOINC has been endorsed by the American Clinical Laboratory Association and the College of American Pathologists. It has been adopted as an alternate test reporting code by large commercial laboratories including Quest, LabCorp, Mayo Medical Laboratories, and MDS Labs; large HMOs including Kaiser Permanente and Aetna;

governmental organizations including the CDC, DOD, VA, and NLM; and has also been adopted by Germany, Switzerland and two Canadian provinces.

- Current draft proposals for HIPAA electronic claim attachment standards are based on LOINC codes.

- LOINC has also been supported in part by funding from NLM, HCFA, DOD, AHCPR (now AHRQ), and the John A. Hartford Foundation. (NLM N01-LM-9-3517).

- The full LOINC database and RELMA -- a program for searching and viewing the LOINC database and mapping local files to LOINC -- are available at no cost at http://www.loinc.org or on a CD which can be requested via an e-mail message to loinc@regenstrief.org or by a fax request to Standards at 1-317-630-6962.

11.4 Medical Subject Headings (MeSH®)

The reason for referring to MeSH at this point is not motivated by it potential role as a clinical terminology, but rather by its wide spread use for indexing the biomedical literature.

11.4.1 The Thesaurus

MeSH is the National Library of Medicine's controlled vocabulary thesaurus. It consists of sets of terms naming descriptors in a hierarchical structure that permits searching at various levels of specificity.

MeSH descriptors are arranged in both an alphabetic and a hierarchical structure. At the most general level of the hierarchical structure are very broad headings such as "Anatomy" or "Mental Disorders." At more narrow levels are found more specific headings such as "Ankle" and "Conduct Disorder." There are 21,973 descriptors in MeSH. In addition to these headings, there are 132,123 headings called Supplementary Concept Records (formerly Supplementary Chemical Records) within a separate chemical thesaurus. There are also thousands of cross-references that assist in finding the most appropriate MeSH Heading, for example, Vitamin C see Ascorbic Acid. These entries include 23,512 printed see references and 102,346 other entry points.

MeSH Applications

The MeSH thesaurus is used by NLM for indexing articles from 4,600 of the world's leading biomedical journals for the MEDLINE® database. It is also used for other NLM-produced databases that include cataloging of books, documents, and audiovisuals acquired by the Library. Each bibliographic reference is associated with a set of MeSH terms that describe the content of the item. Similarly, search queries use MeSH vocabulary to find items on a desired topic.

MeSH is the source of the descriptors used in NLM's Index Medicus® and is fundamental to the organization of this monthly guide to articles from more than 3,400 international journals.

11.4.2 Establishing and Updating MeSH

The Medical Subject Headings Section staff continually revise and update the MeSH vocabulary. Staff subject specialists are responsible for areas of the health sciences in which they have knowledge and expertise. In addition to receiving suggestions from indexers and others, the staff collect new terms as they appear in the scientific literature or in emerging areas of research; define these terms within the context of existing vocabulary; and recommend their addition to MeSH. Professionals in various disciplines are also consulted regarding broad organizational changes and close coordination is maintained with various specialized vocabularies.

11.5 MED

11.5.1 Introduction

The Medical Entities Dictionary is a large repository of medical concepts that are drawn from a variety of sources either developed or used at the New York Presbyterian Hospital, including the UMLS, ICD9-CM and LOINC. Currently numbering over 75,000, these concepts correspond to coded terms used in systems and applications throughout both medical centers (Columbia-Presbyterian and New York-Cornell). It continues to grow at about 6,000 terms per year, although accelerated growth is anticipated as additional network hospitals are integrated into the NYPH system.

11.5.2 Structure

The terms are brought together in the MED and represented as frames, arranged in a semantic network. Each frame includes information specific to the term, such as its name, its code or codes in various systems, and related textual information (e.g., units of measure for tests, synonyms, etc.). The frames also contain pointers to related terms in the MED's semantic net. Some of these pointers form the multiple hierarchy of the MED, while others provide name-attribute information (for example, the concept "Serum Sodium Test" is linked to the concepts "Serum Specimen" and "Sodium Ion" through the relations "has-specimen" and "substance-measured", respectively).

11.5.3 Integration

The MED has proven to be a powerful tool in two respects. The first is the ability to support multiple applications across the institution. Programs which, for example, display laboratory results do not need to stay synchronized with the various laboratory systems with respect to terminology. Instead, they can obtain test names and units from the the MED for any test term encountered in the clinical repository.

11.5.4 Construction

The second powerful aspect of the MED is to support the knowledge-based editing of itself so that, for example, new test terms can be placed in the appropriate classes for later aggregation (it does this through the use of the semantic links "has-specimen" and "substance-measured", which map individual tests to test classes). A number of knowledge-based tools have been created to support knowledge-based term classification and other maintenance functions, including programs which check for validity and internal consistency.

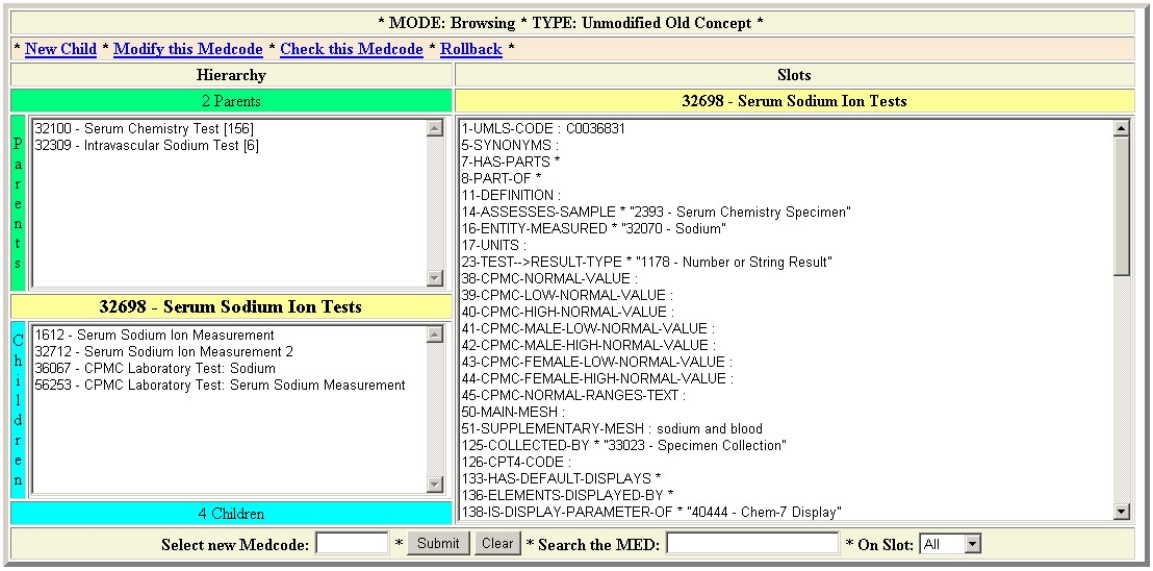

This is a screen shot of the MED Editor, a Web-based program written in C. The current focus of the editor is the class "Serum Sodium Ion Tests" (MED Code 32698). The top-left window shows the two parents above this concept; the bottom-left window shows its 4 children. The window on the right shows the details of the frame-based information about Serum Sodium Ion Tests. Editing functions are carried out through the links at the top of the display.

11.6 MEDCIN

11.6.1 What is MEDCIN?

To be usable, a medical terminology has to describe any clinical presentation in a way that is fast and easy-to-use. It must be in a ready-to-use format with a presentation engine. MEDCIN is a set of clinical terms supplied ready-to-use. The Intelligent Prompting (U.S. patent # 5,823,949) presentation engine presents a short, relevant list of these MEDCIN terms for rapid documentation.

In order to be clinically usable, terms in a medical terminology must be medically related. MEDCIN's presentation engine provides these medical relationships through multiple clinical hierarchies for each MEDCIN term. For example, chest pain is related to fever in the diagnosis of bacterial endocarditis, but not in the diagnosis of angina. MEDCIN includes millions of these medical relationships.

A terminology must provide a computer program with additional information for each term, other than just the words. Each MEDCIN term has an associated property record containing items such as relevant value ranges, units for tests results, laterality flags, control for narrative presentation, cross-references to external code sets, etc

11.6.2 Presenting Relevant Clinical Information

MEDCIN includes more than 215,000 clinical data elements encompassing symptoms, history, physical examination, tests, diagnoses and therapy. Virtually the entire clinical encounter can be encoded. Additionally, MEDCIN's Intelligent Prompting tables provide a means for displaying only clinically relevant items, thus avoiding the problem of overwhelming a user with too much information.

Consider the example of a patient complaining of a cough. Only a very small fraction of the MEDCIN data elements pertain to cough and are relevant to what the user needs. The items needed by the user may vary depending on a patient's age, sex, medical history, and other complaints. Relevant items for a 70-year-old two-pack-a-day smoker should be different from those for a three-month-old child. A cough starting just three days ago, accompanied by a fever, will require different items than for a cough starting nine months ago.

11.6.3 Intelligent Filtering of Medical Information

MEDCIN's diagnostic tables can filter medical information in real-time to present relevant items to the user. This can be done by starting with a symptom and "working forward" (as in the cough example, above) or by starting with a known or suspected diagnosis and getting a diagnosis-specific list. The extent of detail of any such list (short, medium, long) is controlled by the user.

These capabilities are also useful when developing medical applications which require building entry forms or templates (another form of list), either in advance or in real-time. In addition, MEDCIN's diagnostic index can be used to review or filter previously stored data, such as might be used to perform a problem-oriented review of a patient's medical record.

Finally, each MEDCIN finding has a detailed property record, identifying hundreds of attributes useful for automated handling of data. The predictability of these clinical data elements, and the hundreds of associated attributes and relationships, provide the means for building very powerful clinical applications.

11.6.4 Clinically Precise Phrasing

Each MEDCIN data element is a clinical proposition, defined to convey unique intellectual clinical content. For example, consider the following two very similarly worded propositions:

1. "wheezing which is worse during cold weather"
2. "wheezing which is worse with a cold"

Although these two propositions are composed of many of the same words and sub-clinical concepts, they differ significantly to a physician because of their respective clinical implications. These differences in meaning are important to a clinical user and critical for enabling software to present relevant items to a user.

11.6.5 Unique, Permanent and Concept-Independend Identifiers

Each MEDCIN data element has a unique, permanent, non-contextual MEDCIN numeric identification number.

11.6.6 Hierarchical Structure with Inheritance of Clinical Properties Between levels

MEDCIN data elements are organized in a clinical hierarchy. The hierarchy provides inheritance of clinical properties between data elements. This capability greatly enhances the capabilities of clinical software and also provides logical presentation structures for the clinical user and application developer.

11.6.7 Multiple Clinical Hierarchies

Many virtual clinical hierarchies exist in MEDCIN because of the linkage of MEDCIN data elements through their use to describe many diagnoses in the diagnostic index. For example, difficulty breathing is linked to bilateral swelling of the ankles in the diagnostic index of congestive heart failure, but not in the diagnostic index of asthma.

11.6.8 Synonyms

As stated previously, MEDCIN data elements are phrased to convey clinically precise intellectual content. To accommodate differences in expression, MEDCIN provides a table of 600,000+ synonyms and alternate phrasings.

11.6.9 Medcin Narrative Presentations (U.S. patent # 5,802,495)

MEDCIN enables applications to store medical information as coded data elements and produce narrative reports, such as physician encounter documentation, from the same data. The MEDCIN ChartBuilder™ SDK includes an automated narrative generator with a number of different narrative presentation options.

11.6.10 Links to Other Coding Systems

Medcin’s CodeLinks files provide mapping to other coding systems including CPT-4, ICD-9, ICD-10, ICD-O and DSM-IV.

11.6.11 Regular Updates

MEDCIN is updated constantly throughout the year by working with the MEDCIN consulting editors. Updated data files are distributed to licensees two times per year.

11.7 OpenGALEN

11.7.1 Overview: What is GALEN technology ?

Today's Clinical Workstations aren't very clinical: too many systems make the clinician's job harder, not easier. Users complain especially about the medical coding schemes they're forced to use.

GALEN is a radical new technology for medical coding and terminology. Designed as a completely new kind of infrastructure for clinical application builders, it will enable the 'clinical' to be put back into Clinical Workstations.

Healthcare workers expect and deserve clinical workstations that they don't immediately want to throw out of the window. Healthcare managers dream of being able to extract

the information they need from raw clinical information. Healthcare policy makers want ways to get best practice advice onto the desktop, where it is needed and when it is needed.

Meeting these demands requires a new way of representing clinical information within the computer, so that it can work usefully and intelligently with the data it stores. Today's terminologies just don't cut it.

GALEN is a radical new approach to medical terminology for the 21st Century. For the first time applications for end-user data entry, decision support and other data aggregation tasks can use the same terminology without having to use the same terms. The result of over a decade of research and development funded by the European Commission, it is now ready to move into the wider world.

11.7.2 The need for GALEN Technology

The management of clinical information is notoriously difficult

The motivation behind GALEN is to find ways of storing detailed clinical information, about patients, in a computer system so that both:

- clinicians are able to store and review clinically relevant information (at a level of detail relevant to them)
- computers can manipulate what is stored, for retrieval, abstraction, display, comparison, etc.

Traditional ways of storing clinical information in computer systems have difficulties: storing free text (or digitised voice or paper) offers advantages in storage, but the computer cannot manipulate what it holds. Traditional coding and classification schemes are both too big and too small: too big in that it is hard to find elements within them (they are large and not exhaustively classified); too small in that they typically do not provide the level of detail required to support direct patient care.

11.7.3 Need to pull together clinical applications

Many different kinds of clinical applications exist, or are postulated at various levels. For example: decision support, medical records, data entry. The challenge is in finding ways of representing the medical record, say, that can then be successfully interrogated by a decision support system. For example, how does a system determine whether a patient is a candidate for a protocol or guideline?

11.7.4 Need to pull together disparate coding and classification schemes

There are many different coding and classification schemes around and in use; however, they are typically used for only one purpose. It is hard (and often information is lost in one direction or other) to transfer information between schemes. It is hard to re-use schemes for the purpose for which they were originally developed (and this causes the proliferation of even more schemes).

11.7.5 Need to bridge the gulf between classification schemes and clinical applications

We need to bridge the gap between clinical applications, and we need to bridge together different coding and classification schemes. In addition, we need to bring these two worlds together, so that we can, for example, extract abstraction information from clinical records. The "GALEN Bridge" is a visualisation of how we see GALEN fitting in the overall world of coding and classification schemes, and clinical applications.

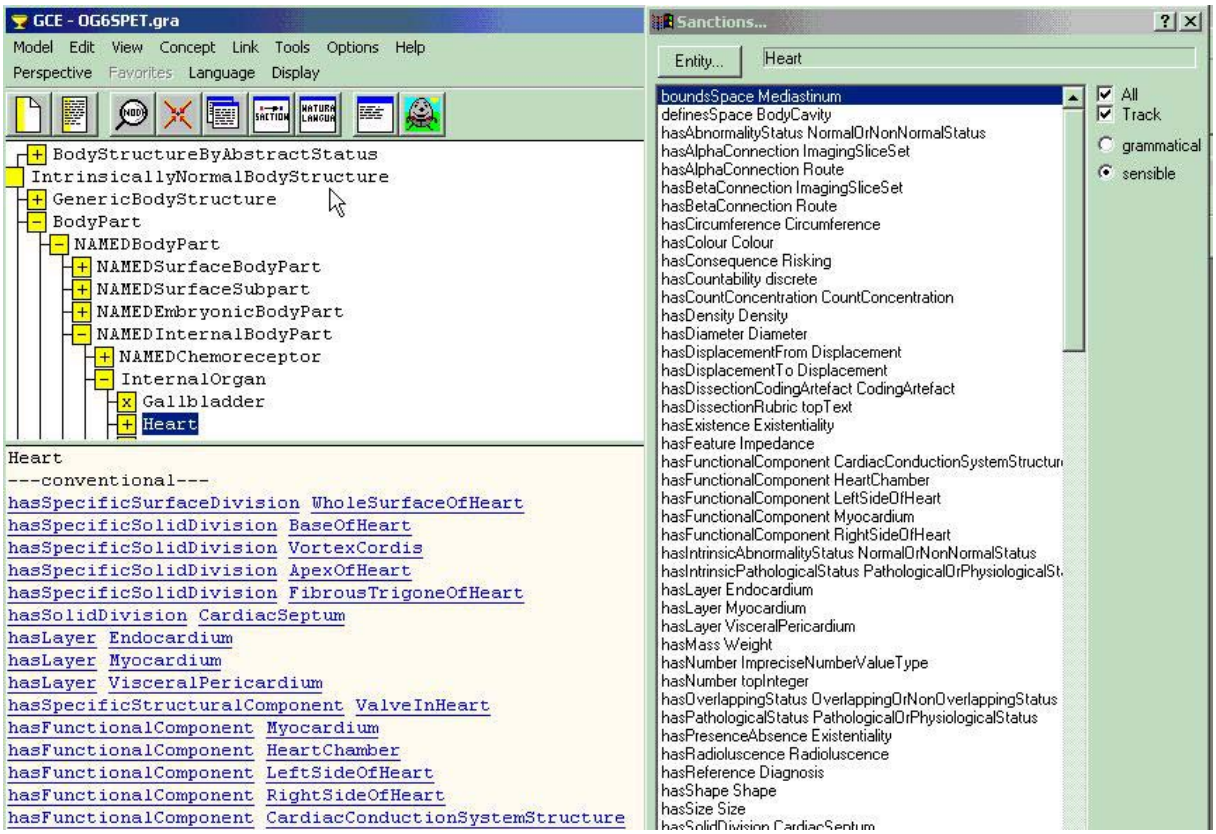

11.7.6 GALEN Functionality

GALEN provides a formal model of clinical terminology

To act as the bridge, GALEN provides a unified representation of clinical terminology: the GALEN Common Reference Model. The model provides an application independent view of clinical terminology, based on a formal representation language (GRAIL) which allows the computer to help n the organisation of the terminology. Rather than putting inheritance links in by hand, for example, the GALEN approach is for the modelling process to be one of saying what clinical concepts are, and then allowing the software to organise the terminology into many coherent, and complete hierarchies.

11.7.7 Terminology Services

Access is provided to the GALEN Common Reference Model through a GALEN Terminology Server: clinical terminologies must be now software rather than just data. Various classes of services are provided:

- Concept services: how can this concept be specialised; what is it a kind of

- Language services: how can this concept be represented in natural language
- Coding services: what is the nearest code in some coding scheme for this concept
- Indexing services: what is the closest information relevant to this concept

11.7.8 Configuration

The GALEN Common Reference Model is application-independent; this, paradoxically, makes it hard to use for a specific application, as it will contain a level of complexity that satisfies the union of many potential applications. Therefore, the server offers different axes of configuration around which it can be specialised to be relevant for a specific application's needs.

11.7.9 Type of Information Managed

11.7.9.1 Clinical information

GALEN manages clinical information, at different levels of abstraction, and in different forms

11.7.9.2 Native representation

For subsequent manipulation, GALEN stores a native representation, either fixed length (which are specific to a particular instantiation of server), or globally shareable, which are not fixed-length (as GALEN can represent arbitrary levels of complexity).

11.7.9.3 Codes

GALEN concepts can be translated to the nearest element of a traditional classification system. In addition, the GALEN Terminology Server provides a 'code store' to store and access existing kinds of coding and classification schemes.

11.7.9.4 Natural language

A concept can be presented as natural language text for display / medico-legal storage in the record.

11.7.9.5 Technical Environment

Implementation platform OS is Win32

Commercial GALEN Terminology Servers are currently available for Win32 (NT/95/98) environments, on top of a proprietary database. Research implementations for other environments (UNIX, MAC-OS) exist. For details of these, contact OpenGALEN.

11.8 Unified Medical Language System (UMLS)

11.8.1 Background

In 1986, the National Library of Medicine (NLM), began a long term research and development project to build a Unified Medical Language System ® (UMLS ®). The purpose of the UMLS is to aid the development of systems that help health professionals and researchers retrieve and integrate electronic biomedical information from a variety of sources and to make it easy for users to link disparate information systems, including

computer-based patient records, bibliographic databases, factual databases, and expert systems. The UMLS project develops "Knowledge Sources" that can be used by a wide variety of applications programs to overcome retrieval problems caused by differences in terminology and the scattering of relevant information across many databases.

11.8.2 Development Strategy

The project is directed by a multi-disciplinary team of NLM staff. NLM encourages broad use of the UMLS products by distributing quarterly editions free-of-charge under a license agreement. The Knowledge Sources are iteratively refined and expanded based on feedback from those applying each successive version. UMLS Knowledge Sources:

There are three UMLS knowledge sources:

- UMLS Metathesaurus ®
- SPECIALIST Lexicon
- UMLS Semantic Network

The Metathesaurus provides a uniform, integrated distribution format from over 100 biomedical vocabularies and classifications (the majority in English and some in multiple languages) and links many different names for the same concepts. The Lexicon contains syntactic information for many terms, component words, and English words, including verbs, that do not appear in the Metathesaurus. The Semantic Network contains information about the types or categories (e.g., "Disease or Syndrome," "Virus") to which all Metathesaurus concepts have been assigned and the permissible relationships among these types (e.g., "Virus" causes "Disease or Syndrome"). NLM also distributes associated lexical programs and software helpful in producing customized versions of the UMLS Metathesaurus. Annual editions of the Metathesaurus have been distributed since 1990.

There have been significant additions and changes to the UMLS Metathesaurus and the SPECIALIST lexicon which should improve their usefulness for many applications. For example, the normalization process used by norm has been changed to allow multiple normalized forms when an input form could be an inflection of more than one base form. This scheme allows a more linguistically motivated uninflection algorithm.

The "MetamorphoSys" software has been improved and is useful in producing customized versions of the Metathesaurus. "MetamorphoSys" facilitates exclusion of any vocabulary for which an additional license arrangement has not been negotiated, or to exclude vocabularies inappropriate to the purposes of the UMLS user.

11.8.3 UMLS Applications

NLM and many other institutions are applying the UMLS Knowledge Sources in a wide variety of Applications including patient data creation, curriculum analysis, natural language processing, and information retrieval. NLM's own applications include PubMed® , the NLM Gateway , ClinicalTrials.gov , and the Indexing Initiative. An issue of NLM's Current Bibliographies in Medicine series, Unified Medical Language System® (UMLS®) covers the structure and semantics of the UMLS Knowledge Sources, their development and maintenance, and assessments of their coverage and utility for particular purposes, and the full range of UMLS applications. It contains 280 citations covering the period

from January 1986 through December 1996. More recent references can be found by searching for Unified Medical Language System on MEDLINE® /PubMed® .

11.8.4 Obtaining the Knowledge Sources

NLM does not charge for the UMLS products. They are available to U.S. and international users. Requestors must sign and submit a License Agreement for the Use of UMLS Products. Licensees are responsible for complying with the restrictions on use of the contents of the UMLS Metathesaurus detailed in the agreement. Some uses of some vocabularies contained in the Metathesaurus require separate agreements with the producers of the individual vocabularies. All the Knowledge Sources are available to licensed UMLS users via the Internet from the UMLS Knowledge Source Server which has a Web interface and an applications program interface (API). ASCII relational files by ftp are available from the Server and may also be requested on CD-ROM. A complete description of the Knowledge Sources and their distribution formats can be found in the UMLS Documentation.