N. F. de Keizer[1, 3]**,
A. Abu-Hanna**[1]**, R. Cornet**[1]**,
J. H. M. Zwetsloot-Schonk**[2]**,
C. P. Stoutenbeek**[3]

[1]Department of Medical Informatics,
Academic Medical Center,
Amsterdam, The Netherlands
[2]Julius Center for Patient Oriented
Research, Utrecht University Medical
School Utrecht, The Netherlands
[3]In cooperation with the NICE foundation, Department of Intensive Care,
Academic Medical Center,
Amsterdam, The Netherlands

Analysis and Design of an Ontology for Intensive Care Diagnoses

Abstract: Information about the patient's health status and about medical problems in general, play an important role in stratifying a patient population for quality assurance of intensive care. A terminological system which supports both the description of health problems for daily care practice and the aggregation of diagnostic information for evaluative research, is desirable for description of the patient population. This study describes the engineering of an ontology that facilitates a terminological system for intensive care diagnoses. We analyzed the criteria for such an ontology and evaluated existing terminological systems according to these criteria. The analysis shows that none of the existing terminological systems completely satisfies all our criteria. We describe choices regarding design, content and representation of a new ontology on which an adequate terminological system is based. The proposed ontology is characterized by the explicit and formal representation of the domain model, the metaspecification of its concepts, the vocabulary to define concepts and the nomenclature to support the composition of new concepts.

Keywords: Ontology, Conceptual Modeling, Terminological System, Nomenclature, Intensive Care, Diagnosis

1. Introduction

As a consequence of the aging population and the continuous development of new medical technologies, there is an increasing volume and intensity of medical services. Due to limited resources, issues such as budgeting, quality assessment and benchmarking [1] are introduced to control costs and quality of care [2]. In intensive care medicine, where costs are high due to expensive technologies and highly specialized personnel, assessment and accountability are particularly important, and are carried out worldwide [3-5]. In the Netherlands, an institution for National Intensive Care Evaluation (NICE) was founded in 1996 [6]. The objective of NICE is quality assessment and quality improvement of intensive care units (ICUs) by comparing case mix (the composition of the ICU patient population) adjusted outcome in various ICUs.

Case mix adjusted outcome means the correction of outcome for variation in population characteristics (corrected for age, diagnoses, severity of illness, etc.). Initially, only the outcome measures mortality and length of stay (the latter used as a crude indication of cost) are considered in this quality assessment, because survival of critically ill patients is the primary concern in intensive care and the chosen outcome measures are objective and easy to measure. In a later stage other outcome measures, including quality of life and costs of care, will be considered as well.

The use of mortality as an outcome measure in quality assessment of intensive care is also found in the prognostic scoring models, such as APACHE II, III [7, 8], SAPS II [9], $MPM_{0/24}II$ [10], which predict case mix adjusted hospital mortality for intensive care patients. These models are mainly based on the severity of disturbance of physiological

parameters measured in the first 24 hours of admission. The APACHE II and III systems are the only systems that also include the primary reason for intensive care admission to adjust for case mix differences. However, the diagnostic categories included in these models have not been empirically derived and are of varying levels of detail. We believe that the use of more structured diagnostic information will improve case mix adjustment and stratification. The confirmation for the importance of diagnostic information in stratifying patients can be found in well-known general medical case-mix adjustment approaches such as Disease Staging [11] and Diagnosis Related Groups (DRGs) [12]. These case mix adjustment methods are mainly based on the underlying health problem or reason for admission. The need for further diagnostic information will increase even more when quality of life is to be assessed.

To date, there has been hardly any or no satisfactory registration of diagnoses in intensive care. In the Netherlands, as in many other countries, intensive care is not a recognized specialty and therefore it was not mandatory to record intensive care diagnoses for national registrations. Intensivists do record, however, diagnoses in free text during the care process. With the introduction of Patient Data Management Systems (PDMSs, which are computer-based information systems that facilitate the collection, integration, retrieval and interpretation of multi-source, multivariant data found in an ICU [13]), a structured and systematic registration of the health status became feasible. This allows to evaluate the patient's health problems over time and correlate these with therapeutic actions. Diagnostic information will only have sufficient quality and is only usable for assessment studies if the registration of diagnoses is part of daily care practice [14, 15], e.g., the recorded diagnostic information in the PDMS is used during shifts.

This paper describes the development of an ontology that facilitates a terminological system for intensive care diagnoses. A terminological system is a system that, based on a specification of concepts and their interrelationship, provides terms that denote these concepts. An ontology formally describes this specification of concepts, relationships and functions in a domain of interest [16]. Hence, the notion of ontology plays an important role in this study. A classification is a terminological system which includes generic relationships between concepts of the terminological system. Although in the literature "classification" is often commonly used for all types of terminological systems, we tried to be specific in terminology about terminological systems. In this paper we use terminology which conforms to the European pre-standard ENV 12264 [17] which describes, among others, the notion of object, concept, attribute, term and code. Briefly, according to the pre-standard, reality can be conceived as consisting of objects. When we think about an object we use properties (information about characteristics of an object) to form cognitive constructs called concepts. Concepts can be hierarchically or non-hierarchically related to each other. Hierarchical relationships can be distinguished in generic relationships and partitive relationships. A generic relationship is denoted by the "Is_a" relationship between a subordinated concept and a superordinated concept, e.g., a mitral valve (subordinated concept) is a heart valve (superordinated concept). A partitive relationship is denoted by the "part_of" relationship, e.g., a heart valve (subordinated concept) is part of the heart (superordinated concept). An object is an instance of a concept, for example a specific mitral valve (e.g., that of patient x) is an instance of the concept mitral valve.

Attributes describe the properties of objects or concepts, e.g. size as a property of an organ. Attribute values quantify attributes. For communication, concepts are designated with linguistic labels called terms. Codes (letters, numerals or a combination thereof) can be used to designate concepts for representation in a computerized system.

The first step in the engineering process of an ontology that facilitates a terminological system for intensive care diagnoses consists of an analysis of the objectives and requirements for such an ontology and the evaluation of existing terminological systems on the basis of these requirements (Sections 2, 3). None of the existing terminological systems proved to completely satisfy the requirements. Therefore, we develop a new ontology and describe our choices concerning the design, the content and the representation of this ontology (Sections 4, 5). The result is a formal specification existing of a meta-model, describing concepts, relationships between them and between concepts and terms in general, such as constraints on the number of terms a concept may have; a domain model describing concepts and their relationships in the IC-domain such as "Health problem" and "Etiology"; the IC-domain, which is formed by concepts such as "Meningitis"; their definition in a vocabulary, e.g., meningitis is an infection of the meninges; and a nomenclature to support the composition of new concepts, e.g., meningococcal meningitis is a meningitis caused by a meningococcus.

2. Objectives and Criteria for a Terminological System to Record ICU Diagnoses

2.1 Objectives

To support the registration of diagnoses in intensive care in a structured and non-ambiguous way, an ontology of intensive care is necessary. For our purposes this ontology has to serve two main objectives:
1. Provide a terminology to adequately describe the health status of a patient to support daily care practice;
2. Provide a structure by which health problems of the patient population can be described on different levels of granularity for the purpose of analysis and evaluation of medical and nursing care.

To realize the first objective (support daily care practice), comprehensive terms are necessary to designate medical problems requiring monitoring and treatment and to improve the communication during shifts between physicians and nurses in and outside the ICU.

To enable the use of diagnostic information for analysis and evaluation of care, which is the second objective, the structure of the terminological system should support aggregation on different levels of granularity. A terminological system with a compositional character (a so-called postcoordinated system) provides the possibility of composing new complex concepts by joining more basic concepts, e.g., infection + meninges + meningococcus. In a purely precoordinated classification each concept, no matter how complex, corresponds to a comprehensive term defined beforehand and has a single representational code, e.g., meningococcal meningitis. Precoordinated concepts are more comprehensible by readily available terms chosen by experts, compared to postcoordinated concepts in which syntax and grammar rules are necessary to guarantee sensible composites. In a postcoordinated system the selection of terms may proceed faster because the structure provides different search entries by using the basic concepts, e.g., a meningitis caused by meningococcus can be reached by infection, meninges and meningococcus. Precoordinated terminological systems are usually

harder to use for aggregation purposes; only the hierarchy or structure in which the clinical concepts are arranged can be used to aggregate diagnostic information. In the case that a precoordinated terminological system has semantic relationships between concepts, aggregation and search are made possible, but most precoordinated terminological systems do not represent these relationships.

2.2 Description of Criteria

In many ICUs, Patient Data Management Systems (PDMSs) have been introduced to continuously collect and integrate various data from bedside devices, a hospital information system and manually entered data, and to present these in a clear format. To integrate the registration of diagnostic information into these systems, a terminological system for intensive care diagnoses should be incorporated into the different PDMSs [18].

Most terminological systems include a structure which is based on the hierarchical relations. Formally these structures respect the features of a semi-lattice in the sense that any two different concepts have a most specific ancestor that belongs to the hierarchy. This feature imposes constraints on the connectivity of the graph implied by the hierarchy. Next to this general feature many authors have attempted to identify the essential conceptual features of an ideal terminological system. In the two left columns of Table 1 the criteria for an ideal terminological system mentioned by Cimino et al. [19] and Campbell et al. [20] are described as far as they are relevant to our objectives. Below, we describe these criteria, related to the two objectives of the ontology for intensive care diagnoses.

The ontology must support the intensivists in describing the diagnoses of patients in a structured way. For daily clinical use, the demand for domain completeness of the terminological system is probably the most important; the user must be able to classify a health problem in its full detail, such as "acute anterior myocardial infarction" or "meningococcal meningitis". Attributes and concepts can be used to add extra detail, such as the attribute value "ante-

Table 1 Evaluation of existing diagnostic classifications. +: the coding system satisfies the criterion; +/-: partial satisfaction; -: the system does not satisfy the criterion.

Criteria	Description	ICM	ICD	Read	Snomed	UMLS
Domain completeness	No terminology restrictions in detail, in breadth nor depth	-	-	+	+	+
Attributes	Modify or qualify meaning of a core concept	-	-	+	+/-	+/-
Synonymy	Multiple non-unique terms for a unique concept	-	-	+	+/-	+
Non-ambiguity/non-vagueness	A concept refers to one and exactly one object in the domain	-	-	+	+	+
Non-redundancy	Each object in the domain is represented by only one concept	+	+	+	-	-
Multiple classification	A concept may be assigned to as many concepts as required	+/-	- (+/-)	+/-	+	+
Explicit relationships	Relationships should be explicitly represented	-	-	-	-	+/-
Context-free identifiers	Codes must be unique and non-significant	-	-	+	-	+
Definitions	A concept's meaning should be (concisely) explained	-	+/-	-	-	+/-
Language independent	Translation across natural languages should be possible	-	-	+/-	+/-	+/-
Mapping (to ICD-9-cm)	Cross references to ICD-9-cm/ICD-10 should be possible	+		+	+	+
Syntax/grammar rules	Complex concepts can be composed by joining basic concepts. Rules must guarantee clinical relevance.	N/A	N/A	+/-	-	+/-

rior" could qualify the anatomical concept "myocard", the attribute value "acute" could qualify the health problem "myocardial infarction" and the concept "meningococcus" could qualify the health problem "meningitis". In this research we use domain completeness as the theoretical completeness of the domain , i.e., whether the domain can be extended in breadth (number of direct subordinates or children of a concept) as well as depth (number of levels in the hierarchy) including formation of newly composed concepts. Hence, we do not examine whether a term is really present in the terminology but rather if it could, in principle, be incorporated in it.

The ontology should allow implementations in which the selection of a concept from the terminological system must be possible by free-text searching or by browsing the concept hierarchy. For free-text searching, synonyms (in more than one language) are highly desirable, as individual users will have their own preferred terms to search for concepts in the domain. Non-ambiguity and non-vagueness guarantee respectively that the meaning of a term is unique and that a term refers to a specific object in the domain. For browsing the concept hierarchy, multiple classification is needed to allow the users to choose their own path to find the desired concept. Non-redundancy should prevent the user from finding different concepts for the same clinical object when following different paths.

As described in Section 2.1, post-coordinated classifications support the aggregation of concepts in a natural way. Syntax and grammar rules should guarantee clinical sensibility of com-

posed concepts. Multiple classification and explicit relationships are also important structure features to make aggregation available in a natural way.

The content of the ontology is not static, but will evolve and change in time; concepts will be added, changed, or declared obsolete. Therefore, context-free codes will be used, that is, the code itself does not hint at the concept it represents. The structure and contents of the ontology have to be extensible. For example, when clinical insights evolve, the structure should not restrict the extension or changes in the ontology. A system with a compositional character includes syntax and grammar rules. These rules support maintenance of the system by keeping the contents consistent and minimizing the effort for updating the contents.

Since admission to the ICU is not an isolated event, information such as reason for ICU admission and complications during ICU admission, has to be available during the rest of the hospital admission. Because diagnostic information of the total hospital admission is usually recorded in a hospital information system using ICD-codes, mapping of ICU diagnostic information to ICD-9-cm or ICD-10 is desirable.

In the next section, existing terminological systems will be evaluated with respect to the above-described criteria.

3. Evaluation of Existing Terminological Systems

We restricted our evaluation to general terminological systems or systems with an intensive care domain, the do-

main which we are most concerned with, and which are intended for use in daily care practice. We examined whether the existing classifications ICNARC Coding Method [3], ICD-9-cm/ICD-10 [21, 22], Read version 3 [23], SNOMED [24] and UMLS [25] satisfy our criteria and, hence, whether they can be used for our objectives described in Section 2. We did not include GALEN [26-28], mainly because GALEN was not operational at the moment we started the development of the new ontology. Furthermore, GALEN has broader goals than ours and for practical reasons we need total control over the software in which the ontology for intensive care diagnoses is developed. However, because of the possibilities of GALEN we intend to assess the GALEN software and CORE model for our purposes in the near future.

The ICNARC Coding Method has an intensive care domain and is a computer-based 5-tier hierarchical terminology of "reasons for admission". It has been developed for prognostic use by the Intensive Care National Audit and Research Centre (ICNARC) and is used in British ICUs. The ICD-9-cm/ICD-10, SNOMED and Read are classifications with a general domain. The UMLS is an ontology which uses different terminological systems to facilitate the retrieval and integration of information from multiple machine-readable biomedical information sources. We have studied the literature and (as far as possible) system documentation to get insight into the terminological systems. We note that it is a general problem that the descriptions of the terminological systems in the literature are usually too informal to really understand their structure and meaning. This hinders the communication in the subject field. Table 1 summarizes the results of the evaluation of these terminological systems. The precoordinated structure of the ICNARC Coding Method, just like the structure of ICD-9-cm/ICD-10, limits the use of these classifications for our objectives. The hierarchical structure of the ICD-9-cm/ICD-10 and the ICNARC Coding Method restricts the theoretical domain completeness. Although the ICNARC Coding Method has an appropriate do-

main, i.e., intensive care, the level of detail is not sufficient to be useful in daily care practice of intensive care. Important features like attributes, synonyms, non-ambiguity, multiple classification, explicit relationships and context-free identifiers are lacking or underdeveloped.

In the Read classification, the existence of synonyms supports different users with different views and language usage in intensive care. The Read classification allows concepts to be placed under more than one parent (multiple classification). However, not all concepts are placed under all possible parents. Due to the availability of some attributes, the Read classification is turned into a (defined) postcoordinated system with flexible granularity in detail. The attributes used in the Read classification provide quantification, anatomical location or specification for the concept to which they are linked. The attributes are not consequently used to describe features of concepts in the domain. This restricts the use of these attributes for aggregation. Furthermore, there is lack of a formal specification, which controls consistency between combinations of attributes of one concept.

The merits of SNOMED and UMLS lie in their postcoordinated structure which optimizes the possibilities for extension of the terminology, and searching and aggregation of concepts. However, in both systems clinical sensibility of the composed concepts and grammatical correctness of their terms cannot be guaranteed due to the lack of rules and guidelines for composing new complex concepts. Clinical judgment of the composed concepts or aggregates is still necessary due to the possibility of clinically nonsensical composed concepts.

Since none of the evaluated terminological systems completely satisfies all criteria, we developed a new ontology that incorporates the strengths of the various terminological systems described above. We formally describe this new terminological system in an attempt to overcome the limitations to the informal descriptions in the literature of terminological systems.

4. Design Choices

In this section we clarify the representation formalism used and the design choices made for a new ontology, based on the criteria described in Section 2 and the analysis of the features of the existing terminological systems described in Section 3.

4.1 Representation Formalism

A good conceptualization of a domain is a description which captures with fidelity the underlying concepts in that domain and the relationships between them. We require a representation formalism of our model which (1) is intuitive and conceptual, i.e., it lends itself for human comprehension and communication, (2) has adequate expressive power, and (3) has a formal basis to avoid ambiguity. The Entity-Relationship (ER) formalism and its extended form (EER), which are widely used instruments in the database community, satisfy most of these criteria [29]. ER is an attractive conceptual model that is capable of expressing concepts (entity types), attributes, relationships and constraints (e.g., cardinality constraints on relationships). However, the ER formalism may not always be suitable for expressing complex constraints. Hence, a more expressive formal instrument is needed to provide the semantics of the ER model and to complement it in order to capture complex constraints. This formal specification formalism is chosen to be many-sorted first-order predicate logic (FOL hereafter, for brevity) due to its universality, expressive power and its proven usefulness in representing clinical data [30]. There are, of course, many other candidate formalisms which could have been selected. For example, one could have equivalently used the popular object-oriented (OO) modelling notations such as OMT [31] or UML [32] instead of ER. Note, however, that the OO approaches provide more notational machinery, such as for the functional and dynamic aspects, than is strictly necessary for our purpose here. Similarly, one can consider conceptual graphs [33] or KIF structures in Ontolingua [34] as good alternatives to pure FOL. FOL is, however, more familiar by

publications from various research communities.

4.2 Concepts and Terms

Synonymy, as well as nonredundancy and multilingual support can be realized by means of separating concepts from their terms, that is, their names. This is shown in the ER diagram in Figure 1. The diagram indicates that a term may denote more than one concept and that a concept is described by terms. We would also like to stress that a term can be preferred for only one concept in a language and that each concept has one preferred term and zero or more synonymous terms per language. This sharpened description of the situation is shown in the lower part of Figure 1 by using FOL.

Note that this is a meta-model of the domain where a 'model concept' in Figure 1 represents any concept in the domain model, such as "Health problems" and "Etiology", which are described in Section 5.

4.3 Precoordination versus Postcoordination

As described in Section 2.1, there are two fundamentally different ways for the structuring of a terminological system: precoordination and postcoordination. There is a clear trade-off between the power of expression and the compactness of the terminological system. Precoordination is characterized by (a) sensible, non-ambiguous and non-vague terms with single representational codes; (b) terms which are useful in daily practice; (c) limited power of expression; (d) combinatorial explosion of terms in attempting to cover a domain. Postcoordination is characterized by (a) composing complex concepts by combining more basic ones; (b) grammar rules to control compositions; (c) great power of expression; (d) compactness of representation.

Terms used in precoordinated classifications can be more natural and useful compared to composed terms, e.g., "Meningococcal meningitis" is more comprehensible than the composition of the health problem "Meningitis" with the associated etiology concept
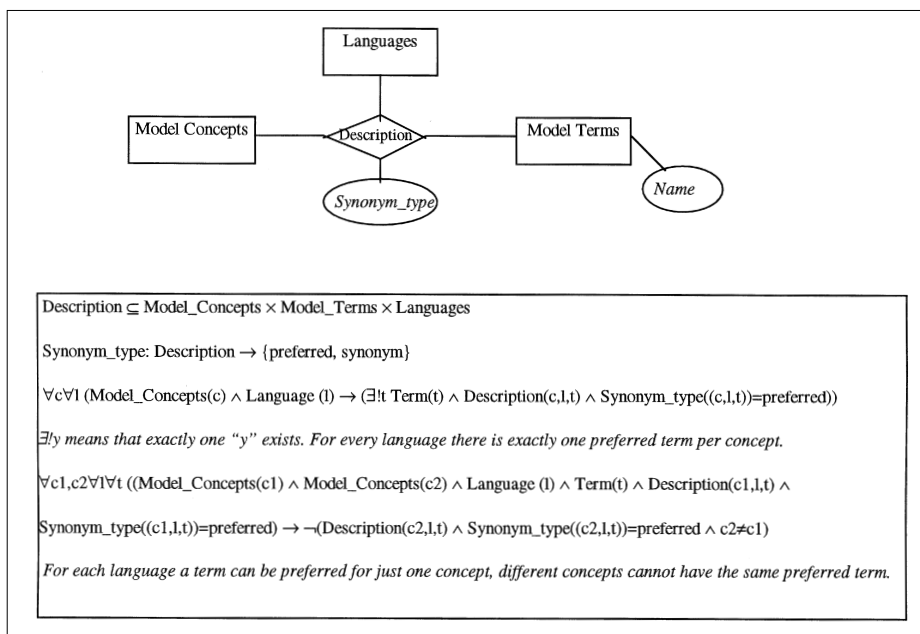

Fig. 1 ER model and FOL description of the relationship between model concepts and terms.

"Meningococcus". The power of expression is closely related to domain completeness. As an attempt to define all possible concepts (precoordination) will lead to combinatorial explosion with respect to the number of concepts, the classification should enable users to compose their own concepts (postcoordination) which makes compact classification possible. A set of syntactical and grammatical rules is necessary to prevent the composition of clinically nonsensical concepts.

To strike a balance between expressiveness and compactness, a combination of precoordination and postcoordination can be sought in which a collection of selected precoordinated terms can be extended, turning it into a postcoordinated terminological system. A mechanism to check whether "new" composed concepts already exist as a precoordinated concept is essential to control consistency within the system, e.g., during the maintenance of the system.

4.4 Clinically Sensible and Non-Ambiguous Concepts

In our model, concepts have to be clinically sensible (nonvague) and non-ambiguous. Some authors seem to summarize these features by the term

holistic concepts. Some health problem concepts can be defined unambiguously and uniquely by associated concepts, e.g., "Meningococcal meningitis" can be described by "nervous system (system) – meninges (anatomical component) – infection (dysfunction) – meningococcus (etiology)". In other cases a health problem concept is ambiguous, i.e., not uniquely defined by the associated concepts. For example, the health problem "SIRS" (systemic inflammatory response syndrome) is a frequently occurring syndrome in intensive care that is characterized by a systemic inflammation and that may or may not be caused by an infection (if it is caused by an infection it is called sepsis). It cannot be localized in a specific body system or anatomical component and has many etiologies. This definition is said to be vague since the term does not stand for a concrete object in the domain. For this reason, the (holistic) term "SIRS" describing this holistic concept is by itself always necessary to define the specific health problem. In other words an intensional definition of a concept (a definition based on a set of characteristics comprised by a concept) is not always possible.

In postcoordinated systems a concept can be extended in detail by adding associated concepts or attributes to

Table 2 Relationship between ICD-10 chapters and concepts in the NICE model.

Chapters ICD-10	System	Anatomical component	Dysfunction / abnormality	Etiology
Infectious and Parasitic Diseases			Infection	Micro-organisms
Neoplasms			Tumor/malignancy Hyperplasia	
Endocrine, Nutritional and Metabolic Diseases	Endocrine system Metabolic system			
Diseases of the Blood and Blood-forming Organs	Blood and blood-forming organs			
Mental Disorders			Psychiatric disorders	
Diseases of the Nervous System	Nervous system			
Diseases of the eye and adnexa		Eye		
Diseases of the ear and mastoid process		Ear		
Diseases of Circulatory System	Circulatory system			
Diseases of Respiratory System	Respiratory system			
Diseases of Digestive System	Digestive system			
Diseases of Genitourinary System	Urogenital system			
Conditions of Pregnancy, Childbirth and the Puerperium			Pregnancy related	
Diseases of Skin and Subcutaneous Tissue	Skin and sub-cutaneous system			
Diseases of Musculoskeletal System and Connective Tissue	Musculoskeletal system			
Congenital anomalies			Congenital deformity	
Certain Conditions Originating in the Perinatal Period	N/A	N/A	N/A	N/A
Symptoms, Signs, and Ill-Defined Conditions			Dysfunctions	
Injury and Poisoning			Trauma, Intoxication, Envenomation	
Classification of Factors Influencing Health Status and Contact With Health Services	N/A	N/A	N/A	N/A
Classification of External Causes of Injury and Poisoning				Accident, Toxic substance

make a new composition. Whether a concept should be defined as one holistic concept or as a composed concept is somewhat arbitrary. We have chosen to define health problem concepts at least at a level of clinical relevance for intensivists and IC nurses or we mark a selected set of health problems, that themselves are not holistic, for compulsory extension to result in clinically relevant concepts. For example "meningitis" is a clinically sensible health problem concept although in most cases extra information about the causing micro-organism is available. However, the concept "fracture" without an anatomical location is not sensible as a health problem (it is a sensible dysfunction concept) because one does not know how to clinically interpret this health problem. This concept must be composed with at least an anatomical component such as "femur fracture", which is a clinically sensible health problem. The fact whether it is an "open" or "closed" fracture can then be further defined by making a new composite concept with an attribute value "open" or "closed".

4.5 Concept Codes

In an ontology in which new concepts can be composed, a unique code, which will be recorded in the PDMS, has to be generated automatically. This code, though context-free, cannot be chosen at random: if the ontology is used at a number of independent sites, the same composed concept should generate the same code to facilitate inter-changeability of information between the sites. To solve this problem, composed concepts get composed codes based on the unique codes of the basic concepts which are context-free. Although we are aware that this code construction may conflict with the context-free criterion in the sense that changes in the model can lead to invalid codes, the code itself will never be used to interpret the meaning of the concept. Instead, the nomenclature, which will be described in Section 5.2, will always be used to provide the underlying codes. In case all sites are connected to a shared terminology server, each composed concept can get a random code which will be stored by the server.

4.6 Choices Related to the Content of the Domain Model and IC-Model

In Section 4.1 we described the formalism of our choice for representing the ontology. The next step is to decide on the content of the ontology, i.e., which concepts, relationships and restrictions are to be represented in the domain model, which can be used to describe intensive care diagnoses, to be called health problems in the remainder of this paper.

The divisions and semantic categories (classes of concepts, extensionally defined) used in the existing terminological systems are mostly based on "anatomy or topography" and "etiology or morphology". The domain model of the new ontology will conform as much as possible to those classes generally accepted. Because of the similar domain, especially the divisions used in the ICNARC coding method are used as a basis for our domain model.

The choice of the semantic categories in the domain model, e.g., "anatomical components", and concepts in the IC-model, e.g., "heart", is refined and appended by discussing these with the domain experts. This discussion was meant to ascertain the concepts and features the domain experts use to select and retrieve diagnoses or diagnostic categories. The acquisition of the IC-model follows the specifications and restrictions in the domain model. Besides the semantic category representing the health problem, we have chosen four other semantic categories:

The semantic category "Systems" contains the body systems involved in the health problems. To support the translation of clinical concepts into ICD-10 concepts, we use those chapters of ICD-10 which concern body systems, such as chapter 10 "Diseases of the respiratory system". This division is also comparable with the second hierarchical level used in the ICNARC Coding Method. In Table 2 the relationship between ICD-10 chapters and our chosen concepts is shown.

The semantic category "Anatomical components" contains all anatomical components associated with health problems. We used the Anatomical Localization Classification from the Handbook Standardization of Classifi-

cation and Definitions in Health Care of the Dutch Classification and Terminology Committee for Health [35] as basis for this category.

The semantic category "Dysfunctions or abnormalities" contains organ failures, traumas and processes, such as infections and hemorrhages. We mainly used the third hierarchical level, "process", of the ICNARC Coding Method and a small number of concepts from the UMLS.

The semantic category "Etiology" contains causes of health problems. It contains causes such as microorganisms causing infections, toxical substances causing intoxication, and causes of trauma.

Health problems caused by other health problems, such as "respiratory failure caused by meningitis" have a different kind of causality relationship than etiology. This type of relationship is situated in a vague area, because of the question whether it is part of the terminological system (as a concept) or whether it is part of the registration where two different concepts are recorded along with the relationship between them. In cases where the relationship is weak, as in the example above, we choose the latter solution, that is that "respiratory failure" caused by "meningitis" will be recorded per patient in the Patient Data Management System as a relationship between two distinct recorded health problems.

The chosen semantic categories serve three purposes:

1. Support the user to select the precise diagnosis for individual patients by providing different search entries to reach a diagnosis. For example "vital hepatitis" can be reached by the anatomical component "liver" or the dysfunction/abnormality "infection";

2. Facilitate evaluation research by supporting aggregation of diagnostic categories on the basis of (combinations of) associated concepts;

3. Support the maintenance of the classification when new health problems are to be added, changed or declared obsolete. For example, the terminological system creates a warning and asks for confirmation when a concept is added while there already exists a concept with the same characteristics.
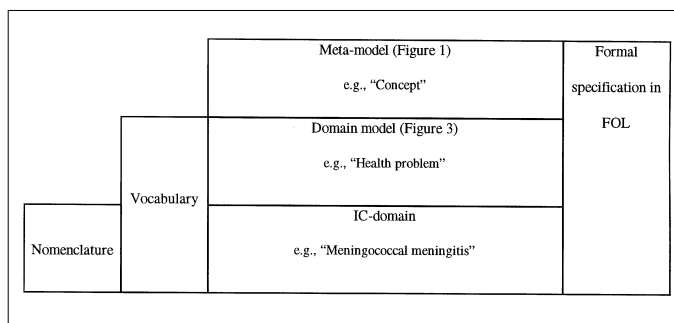
4.7 Composition of Terms and Concepts

The vocabulary and nomenclature, part of our ontology, are formal statements and constraints that describe the health problems and support the composition of new concepts in a way which guarantees that the concepts are clinically sensible. It is possible to specify attribute values and compose new concepts on the basis of the relationships and their specifications defined in the vocabulary and nomenclature, giving the ontology a postcoordinated character. The formal statements and constraints in the vocabulary and nomenclature support the maintenance of the terminological system by checking whether concepts with the same characteristics can potentially be equivalent. Because newly composed concepts are always subordinate concepts generically related to (at least) one parent, which is already present in the nomenclature, the formal statements also support (semi)automated classification.

The relationships between the concepts of the semantic category "Health problems" and the associated concepts from the semantic categories "Systems", "Anatomical components", "Dysfunctions/abnormalities" and "Etiology" can be described by two specifications. Many of the relationships have the specification "define" which means that the health problem concept is (at least partly) defined by an associated concept, e.g., the anatomical component "liver" is part of the definition of hepatitis. These relationships form the vocabulary. The second kind of specification of a relationship is the "specialization" specification which means that extra detail in some given category can be added to further specialize a clinical concept, e.g., possible

microorganisms can be specified for the category "Etiology" as the cause of a health problem. These relationships form the nomenclature.

The way "define" and "specialization" specifications in the vocabulary and nomenclature are used, is described in the following section.

5. Description of the Ontology

The architecture of our ontology is inspired by [36] and consists of six components (Fig. 2): (1) the metamodel and (2) the domain model represented by an ER diagram, (3) their formal specification including complementary information described in first order logic, (4) the IC-domain which is formed by the specialization of the domain model, (5) the vocabulary for a (partial) definition of health problems and (6) the nomenclature for the specification of the allowed terminology compositions.

5.1 The Domain Model and its Formal Description

At the heart of the ER model (see Fig. 3) is the semantic category "Health problems", representing diagnoses in the broad sense of the word. Each health problem may be related to specialized kinds of the health problem, such as "Hepatitis B" is a specialization of "Viral hepatitis" or to subpart health problems, e.g., "Tetralogy of Fallot" consists of the health problems "Ventricular septal defect", "Pulmonic valve stenosis", "Infundibular stenosis", and "Dextroposition of the aorta".

Two different subclasses of health problems can be distinguished: "Direct Health problems" and "OP_Health problems" (monitoring and care after an
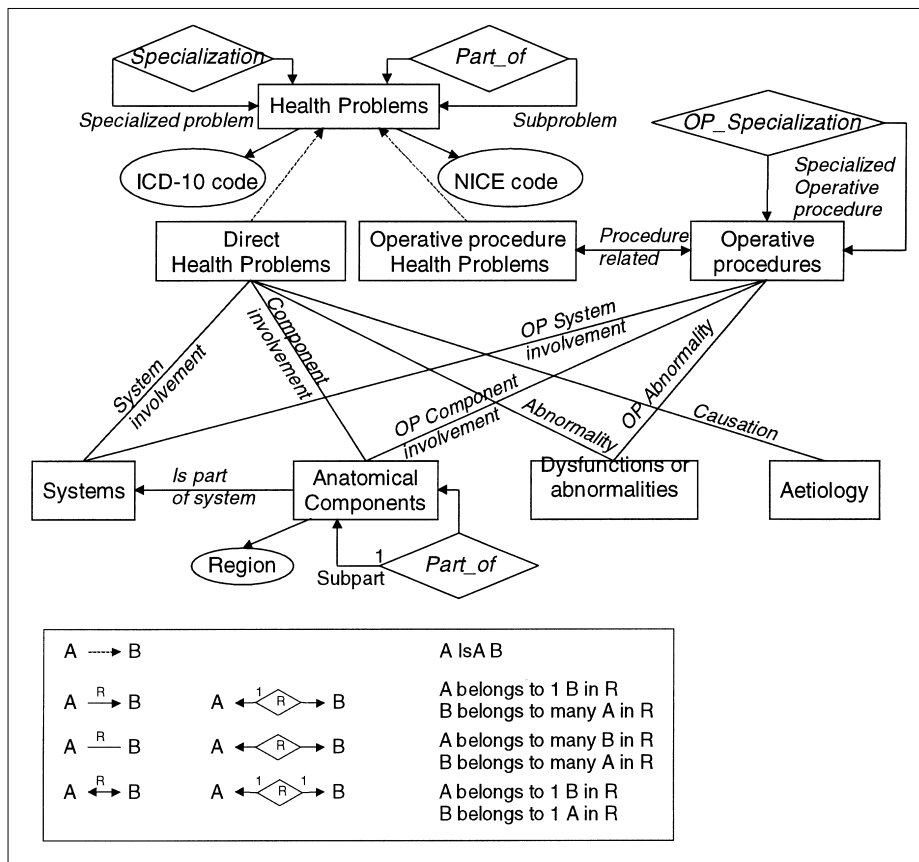
Fig. 3 The conceptual model of the (intensive care) domain model. We used the notations in [41] where numbers indicate cardinality constraints. Note the convention of placing the numbers at the edge of the diamonds, instead of the arrows, in relationship.

ymous, originating from different languages. This relationship is not shown in Figure 3 since it is already specified as a shortcut in the metamodel appearing in Figure 1. The formal specification of the ER model includes additional information that does not appear in the ER model. For example, constraints concerning the description of specialized health problems are formally described in FOL. Asymmetry, for example, is described in FOL by ("$\forall 1, h2_\epsilon$ Health problem Specialized_problem(h1,h2) $\rightarrow -$ Specialized_problem(h2,h1)), which means that when a health problem "h1" is a specialized health problem of the health problem "h2", this health problem "h2" cannot be a specialized health problem of health problem "h1". This constraint is to be inforced by the acquisition tool of the concepts in the first place but cannot be easily described in the ER model.

5.2 The Vocabulary and Nomenclature

The last components of our model are the vocabulary and nomenclature which define the concepts and support the composition of new concepts. In Table 3 a simplified part of the vocabulary and nomenclature is presented. In this example "Direct Health problems" and their related concepts of the four associated semantic categories and the relationship with "Specialized_problem" are presented. "OP_Health problems" are not presented for sake of brevity.

The first row of Table 3 shows that the concept "Viral hepatitis" of the semantic category "Direct Health problem" is defined by "Digestive system", "Liver" and "Infection". This implies that it is not useful to combine the term "Viral hepatitis" with "Liver", since "Liver" is already part of the definition. Furthermore, a specialize-specification, "OR_specialization", exists between "Viral hepatitis" and several viruses mentioned in the semantic category "Etiology". As indicated below, in Table 3, "OR_specialization" prescribes that the concepts mentioned under "Etiology" or one or more of their subordinated concepts can be chosen to further specialize the concept "Viral hepatitis", e.g., theoretically a health problem can be composed of

Operative Procedure). The semantic category "Direct Health problems" is associated with the four semantic categories discussed earlier: "Systems", "Anatomical components", "Dysfunctions or abnormalities", "Etiology". With the exception of "Systems", all semantic categories are top nodes of sub-classifications, in which different levels of detail can be found (this is not shown in Fig. 3 for sake of brevity). This supports the user in selecting and retrieving diagnoses from different views and different levels of detail.

In ICUs, a large group of patients is to be monitored after an extensive operative procedure. For this reason the semantic categories "OP_Health problem" and "Operative procedures" are parts of the conceptual model. This is an example where the domain (intensive care) affects the choice of the concepts. The semantic category "Operative procedures" is linked to the semantic categories "Systems", "Anatomical components", "Dysfunctions or ab-

normalities". If a patient arrives for care after an operative procedure, it is the operative procedure and not the health problem which is related to the associated semantic categories. For example, for a patient arriving at intensive care after a Coronary Artery Bypass Graft (CABG), the health problem "monitoring after CABG" is only linked to the operative procedure "CABG", which is related to the "cardiovascular system", the anatomical component "coronary arteries" and to the dysfunction "obstruction".

Currently we restrict the number of associated semantic categories to the five ones mentioned above, but in the future an extension of semantic categories, which is one of the principles for the design of an ontology [16], is possible. This is contingent on the use of the ontology in daily practice.

Recall that there is a relationship between every concept and its associated terms. Each concept is described by one or more terms, preferred or synon-

two causing microorganisms "Viral hepatitis caused by Cytomegalo virus and Epstein-Barr virus". The concept "Hepatitis viruses" under "Etiology", is related to subordinated concepts, e.g., hepatitis A virus, hepatitis B virus, etc. (Note: although the specification is called "OR_specialization", because it allows one or more concepts to be considered, the result is a conjunction of the concepts chosen instead of a disjunction). "XOR_specialization" is used to indicate that only one of the concepts or one of their subordinated concepts may be used to specialize the health problem, e.g., "Viral hepatitis caused by Epstein-Barr virus". If we should indicate DOR_specialization (Hepatitis viruses), an OR_specialization with a defining component, we restrict the specialization of "Viral hepatitis" to one or more of the subordinated concepts of "Hepatitis viruses" (the concept "Hepatitis viruses" by itself may not be used, as it is a defining property). In accordance to the semantics of "XOR_specialization", DXOR (Hepatitis viruses) restricts the specialization to only one of the subordinated concepts of "Hepatitis viruses".

The concept "Hepatitis B" under "Specialized_problem" in the first row indicates that "Hepatitis B" is a specialized concept of "Viral hepatitis". The second row in Table 3 is the representation of "Hepatitis B". All associated concepts of "Hepatitis B" define this health problem.

In some cases we decided to include complex composed health problems as distinct concepts in the vocabulary and nomenclature, although they could also be composed from a health problem on a higher aggregation level (as "Hepatitis B" can be composed from "Viral hepatitis" in Table 3). The decision to do this is based on the occurrence frequency of the health problem in the intensive care population or on the fact whether intensivists use different terms for the complex composed health problem than can be derived from the composition of the terms describing the composed concepts. Rules will be defined to guarantee similar codes for the same concept, no matter if it is composed or not (Section 4.5). The knowledge engineer is responsible for this process.

Table 3 Part of the vocabulary and nomenclature (simplified).

Direct Health problem	Systems	Anatomical component	Dysfunction/ Abnormality	Etiology	Specialized problem
Viral hepatitis	Def(Digestive system)	Def(Liver)	Def(Infection)	OR_specialisation(Hepatitis viruses, Epstein-Barr virus, Cytomegalo virus)	Hepatitis B
Hepatitis B	Def(Digestive system)	Def(Liver)	Def(Infection)	Def(Hepatitis B virus)	

Def (X): X is an intrinsic characteristic of a concept
OR_specialization(X): a new concept can be composed using X or one or more of X's subordinated concepts
XOR_specialization(X): a new concept can be composed using X or exactly one of X's subordinated concepts
DOR_specialization(X): X is an intrinsic characteristic of a concept, a new concept can be composed using one or more of X's subordinated concepts
DXOR_specialization(X): X is an intrinsic characteristic of a concept, a new concept can be composed using exactly one of X's subordinated concepts

Changes in the vocabulary and nomenclature are not expected to occur frequently but they may occur due to either new medical insights or to modeling decisions of the designer, especially during the phase of setting up the vocabulary and nomenclature. Changes in the vocabulary and nomenclature are local in nature. A typical change in a concept usually implies looking at its immediate parents and immediate children and performing simple checks on them resulting in the modification of their entries in the vocabulary and nomenclature. In this way a change in a concept (such as its move within the hierarchy) does not have to propagate beyond its direct parents and children. This is because of the way concepts are defined in the vocabulary and the modularity of the model, which helps control the effects of changes.

5.3 Implementation

The ontology is implemented in a JAVA application. We have chosen to develop the software in JAVA because of its platform independence. This is an important feature because ICUs use various Patient Data Management Systems running on different platforms. In the future the diagnostic classification has to be incorporated into these PDMSs.

At this moment approximately 1,600 health problems are defined. All health problems and their relationships are defined with the help of domain experts. The JAVA application supports the consistency checking of the domain model and the visualization of the concepts. Synonymous terms still have to be defined. In a pilot study a number of intensivists will evaluate the health

problems and their relationships to the associated concepts defined in the vocabulary and nomenclature for the use of selecting and aggregating diagnoses and will define synonymous terms. We will also use terms used in MeSH and UMLS to make the collection of synonyms more complete.

6. Summary and Discussion

In this section we summarize the objectives of our study, the analysis of the requirements of a terminological system for intensive care diagnoses and the design of a new ontology that facilitates such a terminological system. Finally, we provide some critical observations about existing terminological systems and about our design choices.

Information about the patient's health status and the actual medical problems play an important role in stratifying the patient population for quality assurance of intensive care. To date, no systematic registration of diagnoses, health status and medical problems has been used in daily practice in intensive care, mainly due to the lack of appropriate terminological systems. A terminological system which supports the description of the patient's health problems as part of daily care practice and which supports the aggregation of diagnostic information, is essential. These two objectives resulted in the prioritisation of criteria for terminological systems which were defined by Cimino et al. and Campbell et al. Because none of the evaluated terminological systems completely satisfies all our objectives we decided to design a new ontology using the strengths of the evaluated systems.

We have engineered an ontology for intensive care diagnoses, but the engineering approach and the general ideas are applicable to a broader spectrum of medical domains. We used ER representation techniques to design the metamodel and domain model of this ontology. Because the ER modeling technique is not appropriate to describe all information about the domain we also described the requirements formally in FOL, which is more expressive. However, other equally conceptual and expressive formalisms, such as UML and OMT (for conceptual modeling) and conceptual graphs and Ontolingua (for the formal aspects), could have been used as well. Whatever specific notations are used, the combination of conceptual and formal specifications provides great advantages, which the great majority of current terminological systems do not enjoy, mainly due to the lack of formal specifications. A challenging research issue is the automatic reasoning with (part of) these specifications, for example, as attempted in ontology-based knowledge acquisition such as in GAMES [37] and PROTEGE [38].

In the domain model, semantic categories, which we believe to be useful in describing health problems in the intensive care population, were derived from existing terminological systems. Although the UMLS contains numerous concepts (called semantic types in UMLS), only a very limited number of concepts were useful for our purposes.

There is a similarity between our model and SNOMED. Since we believe that not all health problems can be defined intensionally and since clinical relevance, nonambiguity and nonredundancy are important requirements, we did not follow the postcoordination and compositional terms used in SNOMED, because they lack syntax and grammar rules. There is, however, some research to incorporate semantic information in SNOMED using conceptual graphs, to constrain postcoordination [39, 40].

The Read classification and our model also show some similarities. The main difference lies in the structure: The Read classification uses a strictly hierarchical list of concepts, which can be modified by attributes in a controlled way, whereas we use a conceptual model that includes explicit semantic relationships with explicit constraints. The attributes in the Read code are not consequently used for all concepts in the domain, which restricts the use of these attributes for aggregation, e.g., although the attribute "site" exists, not every concept related to a site is actually linked with a site. Furthermore, formal specifications, which could control consistency between combinations of attributes of one concept, are lacking.

We are conscious of a similarity with the GALEN project [26-28], which is an ambitious project aiming at a formal description of the total medical domain. We did not include GALEN in our current evaluation of existing terminological systems mainly because GALEN was not operational at the moment we developed our ontology. The fact that GALEN is very expressive implies that a great effort should go into the syntax and grammatical rules to guarantee sensible composed concepts. We chose for another balance in expressiveness and compactness, also because we are tackling a specific medical domain, i.e., intensive care. In our ontology, domain knowledge is presented in a domain model, and composition and definition rules are included in a nomenclature and vocabulary. The existence of the "specialize"-specification on relationships, and the explicit constraints on the relationships between concepts give more expressiveness than in precoordinated terminological systems, but take far less effort in sensibility checks than used in GALEN. We think there are also some practical disadvantages to the use of GALEN in our set up, such as the extensive training course necessary before the GALEN software can be used and before a part of the CORE model can be implemented and extended. Furthermore, we require total control over the software by which the ontology and terminological system are developed and this could not be realized with the GALEN software at the time of development. However, because of the possibilities described by the GALEN project we intend to cooperate with the GALEN organization and perform an assessment of the CORE model and the GALEN software to get better insight in the strengths and weaknesses of both the GALEN methodology and our methods.

Currently the ontology implemented in a JAVA application is filled with approximately 1,600 health problems. We are developing a software prototype that consults the ontology and we will incorporate it in the patient data management systems. In the test phase the terminology and the usefulness of the domain model to select and retrieve diagnoses will be evaluated. After sufficient testing and adaptation of the ontology it could be implemented in daily care practice. Once this ontology has been integrated into the daily care practice and into the existing PDMSs, a better insight in the patient population will be possible. This information is essential for the quality assurance and improvement of intensive care.

REFERENCES

1. Keenan S, Doig G, Martin C, Inman KWJS. Assessing the efficiency of the admission process to a critical care unit: does the literature allow the use of benchmarking. Intensive Care Med 1997; 23: 574-80.
2. Relman A. Assessment and accountability. The third revolution in medical care. N Engl J Med 1988; 319: 1220-2.
3. ICNARC Case Mix Programme Dataset Specifications. London: ICNARC 1996.
4. Nelson L. Data, data everywhere (editorial comment). Crit Care Med 1997; 25 (8): 1265.
5. Apolone G, Bertolini G, D'Amico R et al. The performance of the SAPS II on a cohort of patients admitted to 99 Italian ICUs: results from GiViTi. Gruppo Italiano per la Valutazione degli interventi in Terapia Intensiva. Int Care Med 1996; 22 (12): 1368-78.
6. DeKeizer N, Bouman R, Joode J. Anatomical Quality Assurance System (in Dutch). Medisch Contact 1999; 54: 276-9.
7. Knaus W, Draper E, Wagner D, Zimmerman J. APACHE II: A severity of disease classification. Crit Care Med 1985; 10: 818-29.

8. Knaus W, Wagner D, Draper E, Zimmerman J, Bergner M, Bastos P. The APACHE III Prognostic System. Risk prediction of hospital mortality for critically ill hospitalized adults. Chest 1991; 100: 1619-36.

9. Le Gall J, Lemeshow S, Saulnier F. A new simplified acute physiology score (SAPS II) based on a European/North American multicenter study. JAMA 1993; 24: 2957-63.

10. Lemeshow S, Teres D, Klar J, Avrunin J, Gehlbach S, Rapoport J. Mortality Probability Models (MPM II) based on an international cohort of intensive care unit patients. JAMA 1993; 20: 2478-86.

11. Gonella J, Hornbrook M, Louis D. Staging of diseases: A case mix measurement. JAMA 1984; 251 (5): 637-44.

12. Fetter R, Skin Y, Freeman J, Averill R, Thompson J. Case Mix Definition by DRG. Medical Care 1980; XVIII.

13. Milholland D. Information systems in critical care: a measure of their effectiveness. In: Greenes R, Peterson H, Protti D, eds. Medinfo. Healthcare Computing & Communication. Canada, 1995: 1068-70.

14. Grobe J. Editorial comments: Informatics: The infrastructure for quality assessment and quality improvement. J AM Med Informatics Assoc 1995; 2: 267-8.

15. Zielstorff R. Capturing and using clinical outcome data: implications for information systems design. J AM Med Informatics Assoc 1995; 2: 191-6.

16. Gruber T. Towards principles for the design of ontologies used for knowledge sharing. Intern J Human-Comput Stud 1995; 43: 907-28.

17. European prestandard. Medical informatics – Categorical structures of systems of concepts – Model for representation of semantics. Brussel: CEN, 1997.

18. de Keizer N, Stoutenbeek C, Hanneman L, de Jonge E. A PDMS evaluation in Dutch intensive care. Intensive Care Med 1998; 24: 167-71.

19. Cimino J, Clayton P, Hripcsak G, Johnson S. Knowledge-based approaches to the maintenance of a large controlled medical terminology. J Am Med Informatics Assoc 1994; 1: 35-50.

20. Campbell J, Carpenter P, Sneiderman C et al. Phase II evaluation of clinical coding schemes: completeness, definitions and clarity. J Am Med Informatics Assoc 1997; 4: 238-51.

21. International Classification of Diseases, manual of the International Statistical Classification of diseases, injuries and causes of death: 9th revision. WHO 1977.

22. International Classification of Diseases, manual of the International Statistical Classification of diseases, injuries and causes of death: 10th revision. WHO 1993.

23. Read J, Sanderson H, Drennan Y. Terming, coding and grouping. In: RA G, ed. Medinfo 95, 1995: 56-9.

24. Rothwell D. SNOMED-Based knowledge representation. Meth Inform Med 1995; 34: 209-13.

25. Lindberg D, Humphreys B, Mc Cray A. The Unified Medical Language System. Meth Inform Med 1993; 34: 281-91.

26. Rector A, Solomon W, Nowlan W, Rush T, Zanstra P, Claassen W. A Terminology Server for medical language and medical information systems. Meth Inform Med 1995; 34: 147-57.

27. Rector A, Glowinski A, Nowlan W, Rossi-Mori A. Medical-concept models and medical records: an approach based on GALEN and PEN & PAD. J Am Med Inform Assoc 1995; 2: 19-35.

28. Rector A, Bechhofer S, Goble C, Horrocks I, Nowlan W, Solomon W. The Grail concept modelling language for medical terminology. Artif Intell 1997; 9: 139-71.

29. Chen P. The Entity-Relationship model; toward a unified view of data. ACM Trans on Database Systems 1966.

30. Campbell K, Das A, Musen M. A logical foundation for representation of clinical data. J Am Med Informatics Assoc 1994; 1: 218-32.

31. Rumbaugh J, et al. Object-Oriented Modeling and Design. Prentice-Hall 1991.

32. Booch G, Rumbaugh J, Jacobson I. Unified Modeling Language User Guide. Addison-Wesley 1998.

33. Sowa J. Conceptual structures. Addison-Wesley 1984. Reading M, ed.

34. Gruber T. A translation approach to portable ontology specifications. Knowledge Acquisition 1993; 5: 199-220.

35. Dutch Classification and Terminology Committee for Health. Handbook Standardisation of Classification and Definitions in Health Care. Zoetermeer: WCC 1990.

36. Abu-Hanna A, Jansweijer W. Modeling application domain knowledge using explicit conceptualization. IEEE-Expert 1994.

37. Heijst G, Falasconi S, Abu-Hanna A, Schreiber A, Stefanelli M. A case study in ontology library construction. Art Intell, Med 1995; 7: 227-55.

38. Tu S, Eriksson H, Gennari J, Shahar Y, Musen M. Ontology-based configuration of problem-solving methods and generation of knowledge acquisition tools: the application of PROTEGE-II to protocol-based desicion support. Art Intell Med 1995; 7: 257-90.

39. Campbell K, Musen M. Representation of Clinical Data Using SNOMED III and conceptual graphs. In: ME F, ed. SCAMC: McGraw Hill, 1992: 354-8.

40. Rothwell D, Coté R. Managing Information with SNOMED: Understanding the model. SCAMC 1996: 80-3.

41. Wieringa R. Requirements engineering. Framework for understanding. John Wiley & Sons, 1996.

Address of the authors:
Nicolette F. de Keizer, M.Sc.,
Department of Medical Informatics, J2-256,
Academic Medical Center,
P.O. Box 22700,
1100 DE Amsterdam,
The Netherlands
E-mail: n.f.keizer@amc.uva.nl