Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine

AL Rector, PE Zanstra, WD Solomon, JE Rogers, R Baud, W Ceusters, W Claassen, J Kirby, J-M Rodrigues, A Rossi Mori, EJ van der Haring, J Wagner

Appeared as

Rector, A. L., Zanstra, P. E., Solomon, W. D., Rogers, J. E., Baud, R., et al. (1999). "Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine." *IEEE Transactions on Information Technology in BioMedicine* **2**(4): 229-242.

Address for correspondence: Alan L. Rector, MD PhD
Institution: Medical Informatics Group, Department of Computer Science, University of Manchester
Address: Oxford Road, Manchester M13 9PL, UK
Telephone: +44-161 275-6188 FAX: +44-275-6204 E-Mail address: rector@cs.man.ac.uk

Affiliations of other authors

JE Rogers, MRCP
WD Solomon BS
as primary author.

Prof Jean-Marie Rodrigues, MD
Institution: University of Etienne
Address: St Jean Bonnefonds, F-42650, France
Telephone: +33 77 42 79 74 **Fax: +**33 77 42 72 60 **Email:** rogrigues@univ-st-etienne.fr

Robert Baud, PhD
Judith Wagner PhD
Institution: University Hospital of Geneva,
Address: 24 Rue Micheli du Crest,Geneve 4, Switzerland
Telephone: +41 22 372 6203 **Fax:** +41 22 372 6255 **Email:**rb@diogenes.hcuge.ch

Werner Ceusters, MD
Institution: Office Line Engineering
Address: Hazenakkerstaat 20, Sint Lievens Houtem, B-9520 Belgium
Telephone: +32 53 62 2457 **Fax:** +32 2 268 7249 **Email:** wc@ole.medconf.be

Angelo Rossi Mori, BS
Institution: Instituto Tecnologie Biomediche Consiglio Nazionale Delle Recerche.
Address: viale Marx 15, Roma, 00100, Italy
Telephone: +396 827 71 01 **Fax**: +396 827 3665 **Email**: angelo@color.irmkant.rm.cnr.it

Pieter Zanstra, PhD
Egbert van der Haring, PhD
Wim Claassen
Institution: Katolieke Universiteit Nijmegen, Faculteit der Mediosche Wetenschappen
Address: P O Box 9101, Nijmegen, NI-6500 HB, Netherlands
Telephone: +31 24 361 5430 **Fax: +**31 24 361 3505 **Email:** P.Zanstra@mie.kun.nl

John Kirby
Institution: Semantic Technologies Limited
Address: Enterprise House, Manchester Science Park, Manchester M13 6SE
Telephone: +44-161-226-5855 **Email:** jkirby@cs.man.ac.uk

Abstract

A common language, or terminology, for representing what clinicians have said and done is an important requirement for individual clinical systems, and it is a prerequisite for integrating disparate applications in a distributed telematic healthcare environment. Formal representations based on description logics or closely related formalisms are increasingly used for representing medical terminologies. GALEN's experience in using one such formalism raises two major issues:

- How to make ontologies based on description logics easy to use and understand for both clinicians and applications developers;
- What features are required of the ontology and description logic if they are to achieve their aims.

Based on our experience we put forward four contentions: two relating to each of these two issues:

- That natural language generation is essential to make a description logic based ontology accessible to users;
- That the description logic based ontology should be treated as an 'assembly language' and accessed via 'Intermediate Representations' oriented to users and 'Perspectives' adapting it to specific applications;
- That independence and re-use are best supported by partitioning the subsumption hierarchy of elementary concepts into orthogonal taxonomies, each of which forms a pure tree in which the branches at each level are disjoint but non-exhaustive sub-concepts of the parent concept;
- That the expressivity of the description logic must include support for transitive relations despite the computational cost, and that this computational cost is acceptable in practice.

The authors argue that these features will be necessary, though by no means sufficient, for the development of any large re-usable ontology for medicine.

1. Introduction

There are an increasing number of efforts to develop formal representations of medical terminology and concepts [1-6]. These efforts are part of a general move to provide a common infrastructure to re-use clinical information and to integrate the disparate applications in the healthcare environment – *e.g.* medical records, decision support, information retrieval, and management information systems [1, 5, 6]. Integration is only possible if there is a common language which can be shared amongst applications for representing what clinical users have said and done. Even without the requirement for integration, more expressive representations are a requirement for many single applications, e.g. [11, 12]. Indeed, one author has cited establishing a 'common controlled vocabulary' for medicine as the number one 'grand challenge' for medical informatics [7].

A number of the prominent efforts to provide more formal representations for medical concepts are based on 'ontologies' expressed in description logics [8, 9] (or the closely related formalism of Conceptual Graphs [10]). Description logics are the formal descendants of frame systems and semantic networks via KL-ONE [13] and its descendants such as CLASSIC[14], LOOM[15], and BACK[16]. The definition of 'ontology' varies amongst authors [17], but we shall take it here to be the high level schemata and content of a knowledge base expressed in a description logic or similar formalism. The description logic defines the formal properties and semantics of the system – *e.g.* the criteria for one concept being classified as kind of another and the logical operations supported, such as negation, transitive closure, etc. The ontology uses the description logic to describe what exists in the domain — *e.g.* diseases, tumours, anatomical parts, processes, etc. However, although these formalisms and the ontologies built with them share many characteristics, there are also differences between them which require choices, both in terms of the features of the description logic and the conventions for constructing the ontology.

The GALEN programme is among the largest of the efforts to develop ontologies for clinical medicine based on formal representations. It is now moving from experiment to practical demonstration and products. The experiences reported in this paper are based on its two most advanced applications: the PEN&PAD data entry system (now marketed as Clinergy™) [18, 19] and the Classification Workbench (ClaW) which is designed to assist in developing and integrating classification systems [20, 21].

This paper focuses on two issues highlighted by GALEN's experience which we believe must be addressed by any effort to build a formal representation of medical terminology:

1. How to make ontologies based on description logics easy for both clinicians and applications developers to use and understand;

2. What features are required of the underlying ontology and description logic if they are to achieve their aims.

This paper makes two contentions concerning each of these two issues.

With respect to making ontologies usable and understandable:

1.1. That natural language generation is essential to make a description logic based ontology of medicine accessible to clinical users;

1.2. That the description logic based ontology should be treated as an 'assembly language' and accessed via 'intermediate representations' oriented to users, and 'perspectives' adapting it to specific applications.

With respect to the properties required of the ontologies and description logic:

2.1. That the use of orthogonal taxonomies for the subsumption of elementary concepts is an effective strategy to ensure that the ontology is re-usable and modularised;

2.2. That the expressivity of the description logic must include support for transitive relations despite the computational cost, and that this computational cost is acceptable in practice.

For lack of space, and because they have been well covered elsewhere or by other authors, this paper does not discuss a range of equally important issues, *e.g.* overall requirements for clinical terminologies [3], the role of semi-formal analysis in the knowledge acquisition process [22-24], the requirements for managing change in vocabularies and formal concept representations [25-27], or for indexing information for decision support [28]. Nor does it discuss the details of the GALEN representation language (GRAIL) or the GALEN ontology, both of which can be found elsewhere [29-31], or recent more detailed work on the expressivity and optimisation of suitable description logics [32-35]. A summary of the GRAIL language constructs relevant to this paper is given in the appendix.

A note on vocabulary. The vocabulary used to describe formal representations varies from author to author, and from community to community. Table 1 gives rough equivalences between the vocabulary used in GALEN and that used in other related formalisms, and by the CEN ENV on Models of Semantics [36]. Particularly note that GALEN uses the word 'attribute' simply as a generic term for a semantic link whereas description logics and object oriented analysis use the word 'attribute' to indicate links with particular properties.

Table 1: Rough equivalencies of vocabulary between related systems.

GRAIL	***Description Logics***	***Conceptual Graphs***	***Frame systems***	***CEN TC251 'MOSE'***	***Object Oriented Analysis***	***Smalltalk/ JAVA/ C++,...***
Category	*Concept*	*Type*	*Class*	*Semantic Category*	*Class*	*Class*
Individual	*Instance*	*Individual*	*Instance/ object*	*Individual*	*Instance*	*Instance*
Attribute	*Role/ Attribute*	*Attribute*	*Slot*	*Semantic link*	*Relation/ attribute*	*Method*

2. The GALEN Experience

2.1 Making ontologies accessible to users

Most potential users of clinical information systems and most domain experts who must compile the knowledge used in them find the notations of formal representations off-putting at best, and incomprehensible at worst. Furthermore, there is a paradox in developing any generalised re-usable software or representation: the more generic and re-usable it is, the more difficult it is to use for any specific application. To be re-usable, the ontology must address the union of the requirements of all potential applications. This union inevitably includes specialised requirements which are irrelevant to most applications but which can only be satisfied by features with global consequences.

If users are not to become lost in the resulting complexities, there are two options: to present them with natural language generated from the formalism, or to present them with an 'intermediate representation' which hides irrelevant complexities. GALEN uses both techniques: natural language generation when the primary goal is to present information to users - *e.g.* to browse a medical record or classification scheme – and intermediate representations when users must interact with the representation – *e.g.* when developing or manipulating classification systems.

2.1.1 Natural language generation and analysis

The traditional means of attaching language to a 'controlled medical vocabulary' is by attaching fixed text 'rubrics' to codes. This approach presupposes that all possible terms can be enumerated. However, increasingly it is clear that no enumerated list of terms can ever be adequate to all applications. A description logic allows users to compose new descriptions from simpler existing concepts as needed - *e.g.* to compose the concept of a 'severe chronic pneumonitis of the left lower lobe caused by phenol inhalation' – a perfectly understandable medical concept but not one found in any common classification system because it may not yet have occurred sufficiently frequently to be included on any fixed list. There are indefinitely many such expressions. Hence no strategy of assigning text to a fixed list of terms can succeed in expressing the full richness of medical expressions.

The alternative is to generate the language expressions from an underlying representation. Of course, the 'natural' language generated from a formal representation is unlikely to be idiomatic. There is an entire sub-field on language generation devoted to solving these problems (*e.g.* [37]). However, GALEN has opted for highly general but in some ways simpler techniques which can be adapted to many different languages [38] [39].

These techniques are designed to produce a noun-phrase from an arbitrarily complex description logic concept; they rely on the annotation of concepts within the ontology with words and other syntactic information such as number and gender, together with additional rules specifying the grammatical form that should be used when representing each part of the underlying structure. Perspectives are used to simplify and transform the concept into a more linguistically-oriented structure, and then a natural language phrase is built up, using information found in the syntactic and grammatical annotations. These techniques are inherently multi-lingual, and have been demonstrated on a number of languages including English, French, Dutch, Italian, Swedish, Finnish, German and Italian.

GALEN finds this approach adequate because a) for purposes of terminology it need only produce isolated noun phases, and b) even stilted phrases in the users' own language give significant added value over either fixed rubrics or formal outlines in a foreign language. For different purposes it can generate different qualities of language generation: either 'pedantic' which is the initial results of the generation and which is useful to see the full complexity of the underlying representation (useful for quality assurance), and 'improved' which is the best attempt to produce a quasi-idiomatic rendering in the target language (used for end-user display).

Figure 1: Simple natural language generated as feed-back from a data entry system based on forms produced from a description logic via 'perspectives' in the PEN&PAD (Clinergy™) system

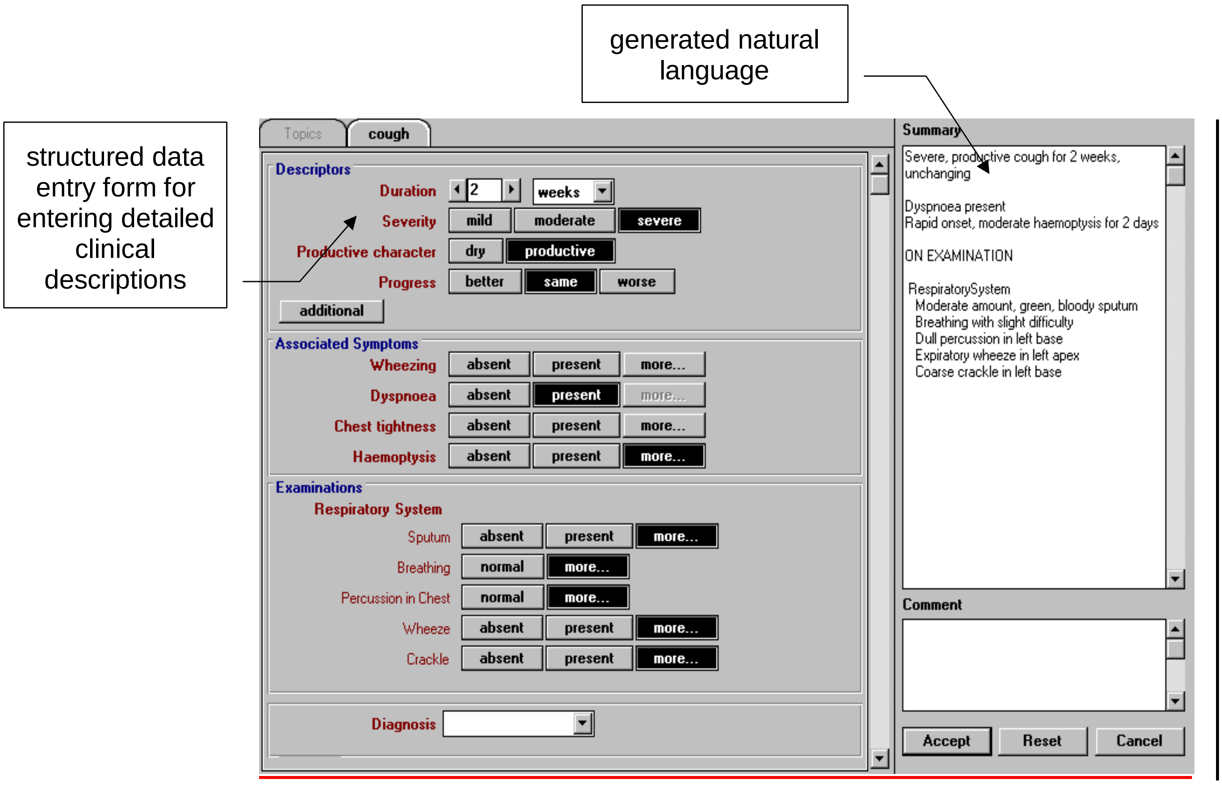

GALEN uses natural language generation in three ways:

1. In the PEN&PAD data entry interface [19, 40]: to produce the primary record seen by the doctors (see Figure 1);

2. In the Classification Workbench for feedback to authors and quality assurance: to let clinical experts authoring or checking terminologies see the results of their work in a familiar form. This is particularly important when merging classifications which originated in different languages. Natural language generation allows the results to be displayed in a single language; without natural language generation there would be no means of any one expert assessing the result. (See Figure 2)

3. In localisation and development of applications: to provide an initial translation into the local language for clinical use or later refinement by local clinicians.

The role of natural language in quality assurance is particularly important even for work within a single language environment. For example, the French government is using GALEN's techniques primarily in order to generate detailed French explanations of the formal representations for quality assurance which are comprehensible to experts[41]. Furthermore, for this purpose, the fact that the language is pedantic and explicit rather than idiomatic is seen as an advantage rather than a drawback.

2.1.2 Intermediate representations and perspectives

Figure 2a: The Classification Workbench (ClaW)

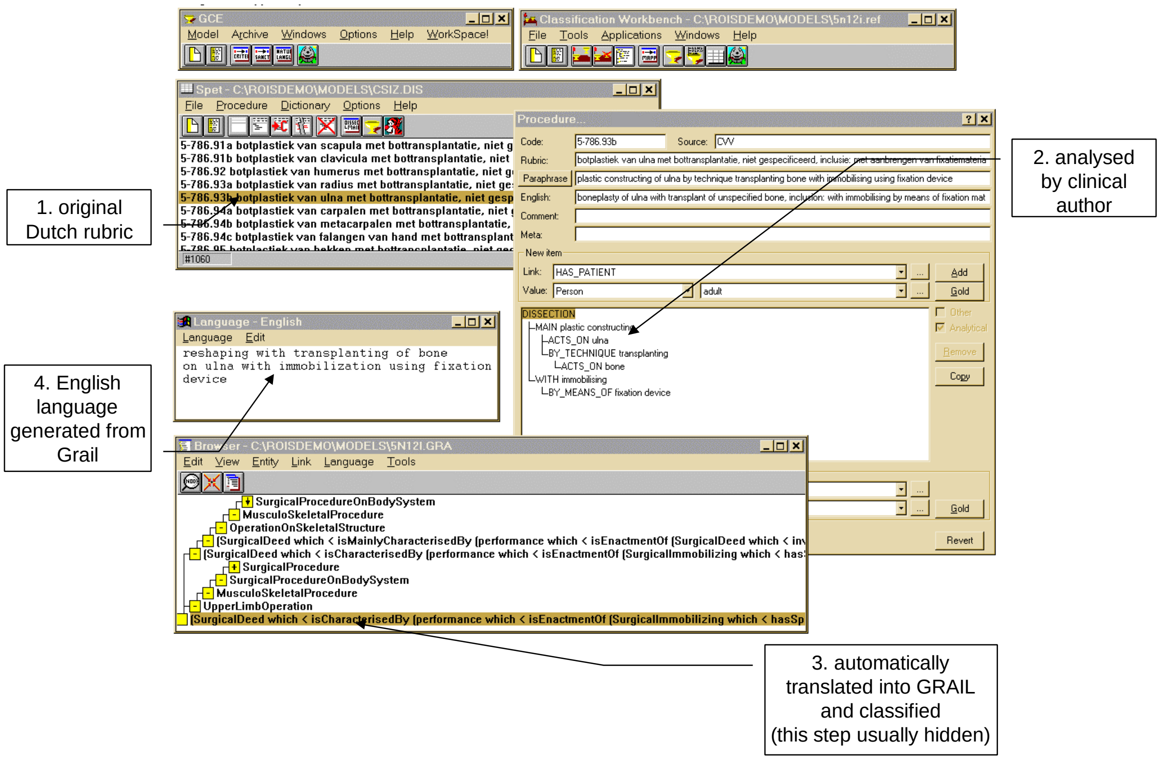

Figure 2b: Detail of content from Figure 2a

RUBRIC	"botplastiek van ulna met bottransplantatie, niet gespecificeerd, inclusie:met aanbrengen van fixatiemateriaal"
PARAPHRASE	"plastic constructing of ulna by technique transplanting bone with immobilising using fixation device"
ENGLISH_RUBRIC	"boneplasty of ulna with transplant of unspecified bone, inclusion: with immobilising by means of fixation material"
SOURCE	"CVV" CODE "5-786.93b"
INTERMEDIATE REPRESENTATION	MAIN plastic constructing ACTS_ON ulna BY_TECHNIQUE transplanting ACTS_ON bone WITH immobilising BY_MEANS_OF fixation device
GRAIL	(SurgicalDeed which < isMainlyCharacterisedBy (performance whichG isEnactmentOf (SurgicalReshapingProcess whichG < LocativeAttribute Ulna hasSpecificSubprocess (SurgicalTransplantation which actsSpecificallyOn Bone)>)) isCharacterisedBy (performance whichG isEnactmentOf (SurgicalImmobilizing which hasSpecificPhysicalMeans FixationDevice))>) hasProjection (('CVV' schemeVersion 'default') code '5-786.93b' 'code')
GENERATED ENGLISH	Reshaping with transplanting of bone on ulna with immobilization using fixation device

Figure 3: Perspectives, Intermediate representations, and applications in the GALEN architecture.

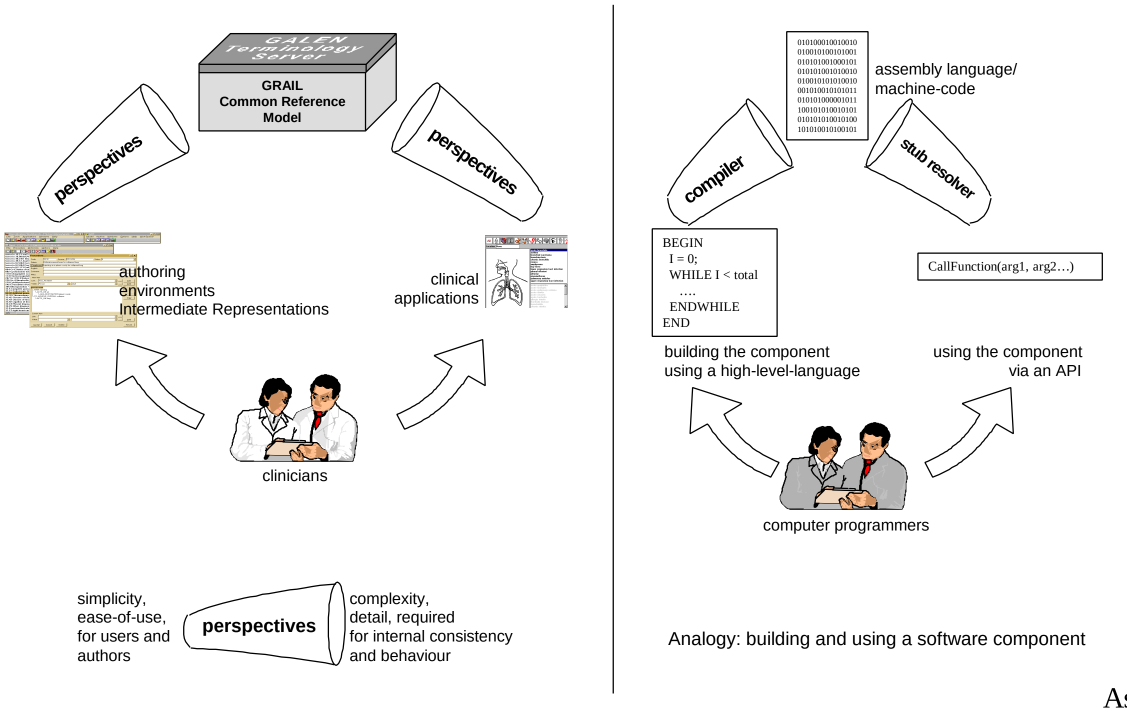

As an example of a specialised requirement with global consequences, consider the following. In order to be able to express the concept of a heart rate which is "elevated, falling, and has a value of 120 beats per minute", as advocated by Shahar [42], GALEN introduced an extra concept *Level* and a corresponding extra degree of nesting in the concept representation:

HeartRate which hasLevel
(Level which <hasQualitativeLevel elevated
hasTrendInLevel falling
hasQuantity 120 beats_per_minute>
)

The concept *Level* is re-usable and can apply to concepts such as temperature, white count, frequency of asthma attacks, etc. However, adding the re-usable intermediate concept of *Level* inevitably makes it more complex to say, simply, 'the heart rate is 120 beats per minute', which the user would naturally like to express as something like:

HeartRate which hasValue 120 beats_per_minute

Insisting on the use of the abstraction *Level* has global consequences which make life more difficult for most users most of the time. Similarly, in order to cope with the rubrics in the procedure coding schemes of the form 'A with B and without C', all procedures are represented internally in a form equivalent to 'performance of A', 'performance of A and non-performance of B', etc. as shown in Figure 2b.

GALEN deals with such cases by treating the underlying representation analogously to an 'assembly language' – *i.e.* an internal representation accessed by specialised external constructs. Applications interact with the description logic via 'perspectives' and expert authors interact with it by means of 'intermediate representations'. The perspectives filter and transform the formal representation so that it matches the particular requirements of the application. The 'intermediate representation' allows the expert authors to enter information in a much simpler, often context dependent format which can be transformed into the description logic representation. The overall architecture is shown in Figure 3 and indicates how perspectives and intermediate representations are used to link clinical expert authors to applications for clinical users via the underlying ontology.

Perspectives have three primary functions in GALEN:

1. Filtering out of unwanted values and unwanted attributes, for example the GALEN ontology requires all structures to have a topology ('solid', 'hollow', 'two-dimensional', etc.) and most phenomena to be marked as to whether they are 'normal' or 'non-normal' and whether they are 'pathological' or 'physiological'. These properties are important for classification and inference internally but irrelevant to most applications and users.

2. Transformations of 'chains' of attributes so that they appear as a single attribute. Taking the example above,

   … hasLevel Level *__which__* *hasQuantity …*

   might be transformed into simply

   … hasValue …

3. transformations which allow more radical re-arrangement of the information, for example to re-arrange

   Pneumonia *__which__* *hasLocation*
   (Lobe *__which__* *isDivisionOf Lung)*

   to something equivalent to:

   Pneumonia
   LOCATION lung
   SUB-LOCATION lobe

   as required by many existing applications and user interfaces.

A major advantage of the use of perspectives is that all of these transformations are outside the core inference mechanism of the GRAIL description logic and therefore do not affect its computational properties. The effect on computational complexity and relative expressiveness of various combinations of features in description logics is far from intuitive (see e.g. [43] [32]). Often, features needed by applications or users would compromise the underlying system if implemented directly, whereas their effects can often be approximated using an external transformation via perspectives. For example, in 2) above, most mechanisms for defining *hasValue* as a new attribute directly within the formalism would make the formalism second order and therefore radically alter its computational properties. (There is a slight loss of expressivity in simulating this effect via perspectives, but this has little effect in practice on most applications.)

Whereas perspectives provide a mechanism for communicating GRAIL constructs to applications, intermediate representations provide an improved mechanism for users, particularly expert clinical authors, to communicate their notions in a form familiar to them but compatible with GRAIL. Intermediate representations are simplified, sometimes context sensitive notations, which can be translated into GRAIL using the perspective mechanisms. Figure 2 shows a typical intermediate representation and the resulting translation into GRAIL.

Intermediate representations have long been advocated by the knowledge acquisition community as a key step in developing knowledge based systems (*e.g.* [44]). Their use in GALEN serves three purposes.

1. The intermediate representations allow clinical experts to author terminologies with a minimum of training because their format and vocabulary are familiar. The vocabulary – or set of 'descriptors'– can be tailored to each site and language and each site's vocabulary then mapped to the common model. The formats – or 'templates' – can avoid most controversial issues, or even allow different centres to use alternative templates which can later be mapped to the same internal constructs.

While adapting the intermediate representation to each site may seem laborious, the gain in local autonomy and ease of use of the local language more than offsets the added effort.

2. The intermediate representations allow changes in the choice of how to represent difficult constructs in the underlying description logic to be made independently of the expert authors. Formal representations are analogous to programming languages. All have limitations, and most offer more than one potential representation for a given notion. Ultimately the techniques to deal with limitations and the choices between alternatives are determined by what works correctly - *i.e.* what representation leads to correct classification. Often, when issues are first encountered, there is insufficient evidence to determine the best technique or choice. For example, in our source material we have found examples of the 'extent' of an operation being described by a large list of terms including 'radical', 'local', 'partial', 'total', 'palliative', 'curative', etc. How best to deal with 'extent' in the formal representation, whether by a single, two or three different attributes required experimentation. However, since in practice there is very rarely more than one value from the set of six selected, users need not be aware of the changes. When it was decided that *palliative/curative* should be split off as a separate attribute, no change was required either to the intermediate representation or to the work already completed by the expert authors.

3. The intermediate representations allow information to be captured from expert authors which cannot currently be represented formally or which requires special treatment outside the description logic proper. For example usages such as 'other', 'not otherwise specified' ('NOS)', and 'not elsewhere classified' are common in existing classification systems such as ICD9/10 [45], SNOMED [46], and the Read Clinical Classification [47]. Rubrics in all these classifications also contain many expressions which are best interpreted as meta information or instructions to coders rather than as part of the literal definition of the concept intended, *e.g.* "cirrhosis *without mention of alcohol*" for non-alcoholic cirrhosis in ICD9. Furthermore, most of these expressions must be interpreted with respect to a particular classification scheme rather than with respect to medicine generally. 'Not elsewhere classified' means 'not elsewhere in *this* classification'. On the one hand, it is important that the expert clinical authors capture all information in a single session. On the other hand, this information is not about the domain of medicine but about the coding system, and hence should not be represented in the core description logic model of medical concepts. The intermediate representation allows all information to be entered together, and then each type of information treated appropriately when processed.

Figure 2 shows an example of the intermediate representation for a surgical procedure, the resulting GRAIL, and the final natural language. The transformations from the intermediate representation to GRAIL include a) the inclusion of the complex 'wrapper' which is used to normalise all procedures; b) mapping of the 'descriptors' from the intermediate representation to the internal concept names used in the GALEN model; c) expansions of the links in the representation to more complex expressions in GRAIL.

In addition, the combined perspective and intermediate representation mechanisms have served two further functions: linkage to standard natural language processing methods and mediation between other applications.

The 'ontologies' used for language processing such as the Penman Upper Model [48] and GALEN's representation differ markedly in many respects. The criteria for the success of the first is its usefulness in language; the criteria for the latter is the correctness of its classification for information retrieval. For example, on the one hand, the use of 'generalised possession' in Penman deliberately conflates distinctions between different kinds of membership and part-whole relation which are vital to GALEN; on the other hand, GALEN often conflates distinctions, such as those concerning the exhaustiveness of a set of options, which are required by Penman. Furthermore, there are many general linguistic patterns which can be expressed directly in most linguistic formalisms, but which have to be transformed before they fit into the GALEN ontology; for example, "right upper lobe of lung" can be expressed directly in most linguistic formalisms but must be transformed to "upper lobe of the right lung" before being represented in GALEN. Ceusters [49] discusses these issues in detail.

GALEN's approach in this situation is to treat the two models as separate and equally valid, and then to develop perspectives to transform from one to the other via specially adapted intermediate representations. This two-stage approach allows us to use standard linguistic techniques with a minimum of modification without undertaking the potentially enormous task of merging the two models with its inevitable risk of compromise to the integrity of each. (However, there is not universal agreement on this technique. Other authors such as Hahn [12] have taken the alternative approach and attempted to merge linguistic and concept models.)

Finally, it is intended that perspectives play an important role in the use of the GALEN concept model for mediation. By linking two different applications via two different perspectives, GALEN can provide for mediation between them across both semantic and syntactic variations, but this mechanism remains to be tested extensively in practical applications. An example due to Rossi Mori [50] is shown in Figure 4 in which the same underlying concept is mapped to two different potential implementations of recent standards for medical records architectures.

Figure 4: Two medical records accessed via different perspectives on the same concept in the Common Reference Model:

Perspective for Record 1	**GRAIL Common Reference Model**	**Perspective for Record 2**
PhysicalExam site Breast Value: value lump laterality right	PhysicalExaminationAct which shows (Lump which hasLocation (Breast which hasLaterality right))	Right Breast Physical exam Lumps Value: present

2.2 Features required of re-usable ontologies and the formal representation

The primary goal for formal representations for re-usable clinical terminologies is that they be capable of being manipulated reliably by computer independent of any specific application. Hence the information to be manipulated must be expressed explicitly in such a way that inferences can be made by general logical methods which do not depend on application-specific processing. Furthermore, if the information is to be maintained, the representation needs to be as modular as possible. Each piece of information needs to be able to be changed independently and the effects of changes should be predictable. The issues are linked: the tacit assumptions behind the organisation of existing terminologies are usually application-specific and are a major barrier to modularity.

2.2.1 Structure of the ontology: explicitness and orthogonal taxonomies

Existing terminologies such as ICD9/10, SNOMED, the Read Clinical Classification, etc. are neither explicit nor modular. They rely on human interpretation of the information expressed in the language of their rubrics plus the human intuitions about the meaning of their hierarchical structure. Significant information is often left completely tacit. Despite efforts by their authors to be more explicit, the language of their rubrics is subject to all the vagaries of human natural language and the general maxim for human communication 'what everyone knows need not be said.' Furthermore, because information on a given topic is often scattered in several sections, numerous special mechanisms are required such as 'excluding', 'including', 'not otherwise specified' or ICD's 'dagger-asterisk' links. Taken together the irregular structure and special mechanisms mean that the effects of a change in one place may have unforeseen consequences elsewhere.

GALEN's approach to making information explicit and modular is based on what we call 'orthogonal taxonomies'. The underlying principles are simple. Any description logic ontology consists of a hierarchy of elementary concepts, *e.g.* 'Heart', 'Lung', 'Disorder', etc. and mechanisms for defining new concepts and/or describing old ones, *e.g.* to define 'Disorder which affects Heart' or "Heart disease"; 'Disorder which affects Lung which is caused by disorder which affects heart', "Cardiogenic pulmonary disease" (e.g. "cor pulmonale"), etc. The central goal is that as much of the information as possible should be contained in the definitions and descriptions where it is amenable to formal manipulation.

Practically this approach is embodied in three principles, that:

P1. The primary hierarchical relationship should be the 'is-kind-of' or 'is subsumed by' relation, strictly understood, *i.e.* to say that "A is-kind-of B" (or equivalently that "A is subsumed by B") is exactly equivalent to saying "all instances of A are instances of B" or in standard logical notation, " $x.\ Ax \rightarrow Bx.$ (This definition of 'subsumption' is central to the semantics of all description logics and related formalisms.)

P2. The elementary concepts should be organised as disjoint non-exhaustive hierarchies , *i.e.,* that each elementary concept should have exactly one elementary 'parent' and be disjoint from its 'siblings' – for example 'Heart' is a child of 'BodyPart' and has siblings 'Lung', 'Liver', etc.; 'Diabetes' and 'Gout' are both children of 'Disease', etc; 'Lung' and 'Liver', 'Gout' and 'Diabetes', etc. do not overlap although they can co-exist. However, it is not generally the case that the elementary 'children' of an elementary concept can be counted on to exhaust that concept. For example, any list of diseases will inevitably be incomplete because new diseases are continually being described, *e.g.* AIDS would not have occurred in any list of diseases before 1980.

P3. All more complex structures and any additional parents should be inferred from descriptions and definitions by the description logic mechanism.

We describe these hierarchies, or taxonomies, of elementary concepts as 'orthogonal' because changes to one taxonomy can be made independently of any changes to any other taxonomy, since no two taxonomies can overlap.

This approach deals with two problems which make existing classifications difficult to use for formal inference and computation:

1. It makes the organising principles explicit and therefore available for formal analysis and verification;
2. It avoids inexplicably tangled hierarchies, and thereby allows the structure to be organised from various alternative viewpoints – e.g. by anatomy, disease process, or physiological system.

As an example of making organising principles explicit, consider the common argument as to whether concepts such as 'wheeze' should be 'symptoms' or 'signs'. GALEN approaches this problem by separating out two orthogonal taxonomies, one for how observations are made, and another for the physical manifestations. Hence we have, as shown in Table 2a, hierarchies of phenomena and social roles which give rise, as shown in Table 2b, to definitions of signs and symptoms.

Table 2a: Orthogonal taxonomies of elementary concepts.

…	…	…	…	…
Phenomenon	Role	Organism	Sex	
Noise	Social Role	Person	male	ObservationMethod
BreathSound	Patient Role		female	Observing
Wheeze	Doctor Role			PlainHearing
				Auscultation

Table 2b: Definitions of Signs and Symptoms in terms of phenomena and roles. Note that each of the notions can be re-used in different contexts -

Doctor = Person which playsSocialRole doctorRole.
Patient = Person which playsSocialRole patientRole
Sign = Phenomenon which isReportedBy Doctor
Symptom = Phenomenon which isReportedBy Patient
Man = Person which hasSex male
Woman = Person which hasSex female
WheezeSign = Wheeze which isReportedBy Doctor
WheezeSymptom = Wheeze which isReportedBy Patient
WheezeAuscultationSign = Wheeze which <isObservedBy Auscultation isReportedBy Doctor>

Table 2c: Resulting hierarchies

Phenomenon	*Person*
Noise	*Man*
BreathSound	*Woman*
Wheeze	*Doctor*
WheezeSign	*Nurse*
WheezeSymptom	
Sign	
WheezeSign	
WheezeAuscultationSign	
Symptom	
WheezeSymptom	

In other words, symptoms are phenomena reported by patients; signs are phenomena reported by doctors. Hence a *wheeze which isReportedBy Doctor* will be classified as a *Sign* and a *wheeze which isReportedBy Patient* as a *Symptom*. The taxonomies of social roles and phenomena can be extended independently, for example to include 'CarerRole' and 'NurseRole'. Other dimensions can be added, *e.g.* to include the sex of the person or the method of observation such as whether the wheeze was heard grossly or only on auscultation, independently of whether it was heard by a doctor, nurse, or carer. The classification of composite terms such as signs and symptoms will be maintained by the software according to the rules of the description logic as shown in Table 2c. Each of the taxonomies can be re-used independently – the physiologic knowledge about ‘wheeze’ is kept separately from the organisational knowledge about doctors, nurses, and carers, and likewise separate from the information about methods of observation. Note that, already in this toy example, the total complexity of the multiple classification challenges what could be maintained manually. Note also, that there is no need to define all possible descriptions in advance. We could compose the concept of a “Phenomenon observed by auscultation” were it required and it would be correctly placed in the hierarchy. Complexity can be deferred until it is required.

Table 3a: Hierarchy of organs showing a mixture of kind-of and part-of

organ	
heart	*(heart* is ***a kind of*** *organ)*
heart valve	*(heart valve* is ***a part of*** *the heart)*
aortic valve	*(aortic valve* is ***a kind of*** *heart valve)*
aortic valve cusp	*(aortic valve cusp* is ***a part of*** *the aortic valve)*

Table 3b: Pure kind-of (subsumption) hierarchy of disorders related to organs shown in Table 3a

disorder of organ
disorder of heart
disorder of heart valve
disorder of aortic valve
disorder of aortic valve cusp

Table 3c: Non-subsumption of alternative construct using same organs shown in Table 3a

surface of heart
surface of heart valve
surface of heart valve cusp

As an example of the problem of tangled hierarchies, consider the hierarchy in Table 3a, slightly idealised from a well known coding system. We have added the annotations making explicit whether the differentiating principle is ‘kind-of’ or ‘part-of’. This type of mixed hierarchy is extremely common in existing classifications and constitutes a major limitation on their use for formal reasoning and their re-usability across applications. For example, the taxonomy in Table 3a works for applications related to disorders because it is indeed true that the corresponding disorders form a true subsumption hierarchy as shown in Table 3b. However, it does not follow that the analogous construction works in all contexts as shown by the example in Table 3c. Although ‘heart valve cusp’ is listed under ‘heart’ in Table 3a, it does not follow from the fact that the ‘surface of the heart’ is covered by pericardium that the ‘surface of the heart valve cusp’ is covered by pericardium. To amplify the point, were the structure in Table 3a to be included in a description logic based ontology, it follows from the definition of the ‘is-kind-of’ relation given in principle 1 above, that it would be equivalent to the statement that “All heart valve cusps are hearts”, or more formally:

" x. HeartValveCusp x → Heart x

This statement is clearly false. Logical systems guarantee only that true conclusions follow from true premises. The effects of including false statements as premises in a logical system are unpredictable.

2.2.2 Features of the description logic: Transitivity and part-whole relations

If false statements such as “all heart valve cusps are hearts” are to be avoided and yet the desired classification obtained, then the description logic must support the basic rule that disorders of parts are disorders of the whole, which is required to go from Table 3a to 3b. More generally, the description logic needs to be expressive enough not to force users to make false statements in order to work around its limitations. Unfortunately, the more expressive the description logic becomes, the greater the computational complexity.

The inference required to go from Table 3a to 3b was not supported by any of the widely used description logics which were available at the outset of the GALEN project – nor, indeed, by any of those currently in widespread use – because such constructs were believed to have unacceptable consequences on the computational properties of the classifier. As a result, GALEN had to design a new representation language – GRAIL – with a general construct for dealing with the relations between transitive relations and subsumption. This mechanism is known as the 'specialised_by' or 'refined _along' schema [29] and is analogous to CycL's TRANSFERS-THRO construct [51] . Formally, the specialised_by axiom schema takes the form for any two relations R_1 and R_2, R_2 transitive,

R_1 *specialised_by* R_2 *implies*
$" x\ y\ z\ .\ xR_1y\ \&\ yR_2z \rightarrow xR_1z$

Although in the worst case, classification using these constructs was known to be intractable, experience quickly showed that performance was acceptable in practice. (Theoretical work has since added support to the belief that this will continue to be so [35, 52].)

GRAIL has found this mechanism for dealing with transitive relations necessary for at least three different functions within the ontology.

1. To allow kind-of and part-of hierarchies to be kept strictly separate in keeping with the strategy of 'orthogonal taxonomies' as shown in the transition between Tables 3a and 3b.

2. To support consistent, indefinitely recursive handling of sublocations, for example, the "draining of fluid from the nucleus of a cyst from the kidney" is automatically classified as an "operation on the kidney" because the location relation is specialised by the part-whole relation. A chain of any length can be managed; by contrast most other representations require a limited set of sublocation 'slots'.

3. To represent syndromes and other nested structures, for example a syndrome, such as 'tetralogy of falot' (a congenital syndrome involving four different defects of the heart and major arteries) is classified as a kind of 'disorder of the heart' using another instance of the same schema used for part-whole relations in 1) above: 'the observation of a syndrome including a disorder' is a kind of 'observation of the disorder'.

3. Discussion

3.1 Making an ontology based on description logic usable

The first two contentions of this paper were that natural language generation and intermediate representations were necessary to bridge the gap between users' and applications builders' requirements on the one hand, and the formal structure of the ontology expressed in a description logic on the other. Unless other researchers are able to provide dramatically simpler ontologies which still achieve the degree of re-use and correctness of classification achieved by GALEN, our experience makes these conclusions almost inescapable. Our experience with the Classification Workbench and PEN&PAD is that the introduction of language generation, intermediate representations, and perspectives are critical steps to user acceptance and practical implementations.

With respect to the Classification Workbench, prior to the introduction of the intermediate representation, material was compiled by experts and then entered by 'druids' centrally, and training to become a 'druid' required three to six months' apprenticeship. Following the introduction of intermediate representations, training for new users was reduced to a few days using workshops. Production also increased dramatically; and there are now more than ten centres involved in compiling the classification. Furthermore, the use of the intermediate representation allowed development to be much more loosely coupled than previously. Individual centres can now make changes – adding new descriptors, even suggesting new constructs – and these can be analysed and processed later. Empowering individual centres has been essential to reaping common benefits from shared development.

Similarly, prior to the introduction of natural language generation, cross comparison of classifications and detailed quality assurance were extremely difficult and time consuming, because users had difficulty reading the output; following the introduction of natural language generation it became immediately possible to see results and correct them. In fact, the French Government seized upon the natural language generation specifically as a means of quality assurance for their new classification of procedures. Completely untrained clinicians were able to examine the generated text and recognise that up to 15% of the submitted candidate rubrics were ambiguous and required revision.

With respect to the PEN&PAD data entry system, without an intermediate layer of 'perspectives', the user would be flooded with all possible choices rather than those which might reasonably be expected to be clinically relevant. Furthermore, without perspectives, the organisation of information would not be as expected. Clinical pragmatics – the order and form in which doctors use information – requires much knowledge other than that present in the simple terminology. The perspective layer provides this extra knowledge.

3.2 Requirements for the ontology and description logic

Our two contentions with regards to the ontology and formalism were that the ontology should be based on orthogonal taxonomies of elementary concepts and that the expressivity of the description logic must be sufficient to support transitivity. Our experience with the first contention is that it has repeatedly proved the key to resolving conflicts between the needs of different applications. Once the taxonomies are properly disentangled and made orthogonal, each application or user can be accommodated by appropriately defined composites such as 'Symptom'. We have yet to encounter a convincing counter-example. Of other authors, those, such as Zweigenbaum [11], who require taxonomies to be strict trees and do not provide multiple classification through descriptions, have focused on specialised applications; those such as Cimino who have allowed uncontrolled multiple hierarchies, have found themselves with hierarchies which are difficult to maintain without transformations analogous to the orthogonal taxonomies advocated here [57].

However, it should be noted that GRAIL is not quite sufficiently powerful to carry through this principle completely in practice. It lacks constructors for the conjunction and disjunction of elementary concepts. As a result high-level elementary concepts such as 'phenomenon' must be inserted as the parent of several different categories to simulate disjunction. Similarly, a few low level values such as 'bilateral' must be inserted to represent conjunctions. However, these cases are rare, documented, and treated as technical shortcomings. The limitation will be removed in future formalisms.

With regards to transitivity, given the importance of anatomy to medical terminology, without a proper mechanism for part-whole relations, mixed hierarchies of part-of and kind-of are inevitable. If, as is commonly done, part-whole hierarchies are simulated by using 'kind-of', then classifications inferred on these deliberately incorrect facts will inevitably be wrong in some cases. There seems little point in having a rigorous logical formalism if its limitations require that the basic data must be falsified.

Likewise, given the complexity of medical terminology, it seems essential that any mechanism intended to be re-usable must not impose arbitrary limitations on the number of levels of sub-locations or subparts. Particularly in the representation of surgical procedures, long chains of sub-locations are common.

GALEN's is not the only solution to the problems of part-whole relations and transitivity, nor does GALEN's solution meet all possible related issues. Other authors have proposed potential mechanisms

[51, 58], [59, 60]. To argue the comparative merits of different schemes is beyond the scope of this paper. GALEN's contention is that some adequate method is necessary.

While this paper contends that dealing with transitive relations is a necessary feature for any formal representation for medical concepts, it certainly does not claim that this is sufficient. Our experience argues for the consideration of a number of further features:

1. Constructs for additional constraints or simple rules equivalent to GRAIL's necessary statements (see [29]);
2. Upper and lower cardinality constraints other than 'one' and 'many', although we do not believe that complete reasoning about upper cardinality constraints, which is known to be intractable, is required – e.g. we doubt that we need to be able to infer from the fact that a person has a brother and a sister that they must have at least two siblings and therefore cannot be classified as a member of a 'two-or-less-child family';
3. True negation - negation must be simulated in GRAIL using constructs such as 'performance'/'nonperformance' or 'presence'/'absence'.
4. Simple conjunction and disjunction of elementary concepts.
5. A construct similar to value-maps or the SAME-AS construct in other description logics.
6. Simple spatial constructs including adjacency and contiguity; for example, to indicate that a 'single' fracture may affect the radius and ulna together.
7. An additional construct related to transitivity of the form 'the skin of the forearm is a part of the skin of the upper extremity', or more generally, 'a part of a division of the whole is a division of the corresponding part of the whole'.

Unfortunately, any of these constructs, alone or in combination, can have serious impacts on the worst case performance of classifiers. There is intensive research currently in progress on optimising description logic classifiers and on their empirical, as well as worst case, performance on real world problems with regular international workshops and congresses (see for example [61]). Future languages will be able to include a wider range of constructs than does GRAIL, but there will still be fundamental choices to be made.

3.3 Integration and Telematics

A clinical ontology is an underlying enabling technology for integration. Perhaps the clearest example of the difference between integration using an ontology and other approaches is offered by the experience of surgical procedures. A single surgical procedure classification across European countries is simply not possible because of differences in both how surgery is practised, and the basis for remuneration. However, by mapping surgical procedure codes from each classification to the GALEN Common Reference Model, we can immediately produce an integrated classification which shows the relationships between the different procedures from different countries, as shown in Figure 5.

Figure 5: An integrated classification which shows the relationships between the different procedures from different countries, also showing a portion displayed as generated English text:

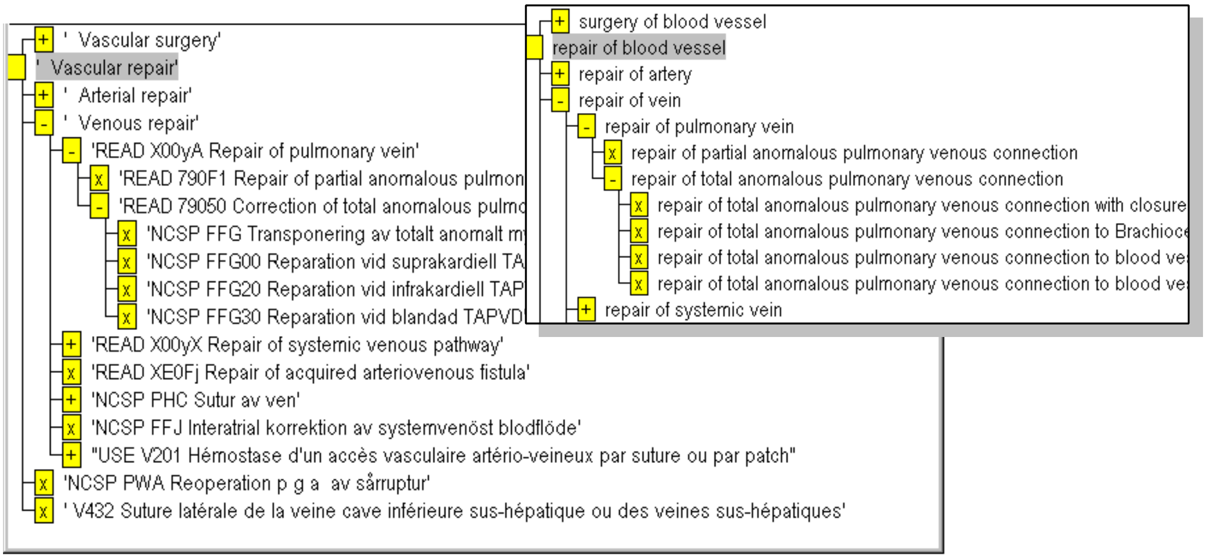

4. Conclusion

The key issues in developing formal ontologies for medicine are to reconcile users' needs with the requirements for formal systems and to establish ontologies which can, in fact, be re-used. GALEN's experience has led to the contention that generated natural language, intermediate representations and perspectives must be used to isolate users and applications from the details of the formal representation, just as high level languages isolate most programmers from the details of assembly and machine language programming. At the same time, it has led to the contention that the ontology should be organised using orthogonal taxonomies of elementary concepts and that the formal representation must provide for rigorous handling of part-whole and other transitive relations. The success of the project in establishing working applications based on large clinical ontologies lends strong support to these contentions. New developments in the underlying technology of description logic and formal classification may change the options available to developers, but we expect these findings to stand.

References

[1] D. A. Evans, J. Cimino, W. R. Hersh, S. M. Huff, D. S. Bell, and The Canon Group, "Position statement: Towards a medical concept representation language," *Journal of the American Medical Informatics Association*, vol. 1, pp. 207-217, 1994.

[2] J. Cimino, "Controlled Medical Vocabulary Construction: Methods from the Canon Group.," *Journal of the American Medical Informatics Association*, vol. 1, pp. 296-197, 1994.

[3] J. Cimino, "Desiderata for controlled medical vocabularies in the twenty-first century," *Methods of Information in Medicine*, pp. (in press), 1997.

[4] K. E. Campbell, A. K. Das, and M. A. Musen, "A logical foundation for representation of clinical data," *JAMIA*, vol. 1, pp. 218-232, 1994.

[5] K. A. Spackman, K. E. Campbell, and R. A. Côté, "SNOMED-RT: A reference Terminology for Health Care," *Journal of the American Medical Informatics Association (JAMIA)*, vol. (Symposium special issue), pp. 640-644, 1997.

[6] A. Rector, P. Zanstra, W. Solomon, and The GALEN Consortium, "GALEN: Terminology Services for Clinical Information Systems," in *Health in the New Communications Age*, M. Laires, M. Ladeira, and J. Christensen, Eds. Amsterdam: IOS Press, 1995, pp. 90-100.

[7] D. F. Sittig, "Grand challenges in medical informatics," *Journal of the American Medical Informatics Association*, vol. 1, pp. 412-413, 1994.

[8] R. MacGregor, "The evolving technology of classification-based knowledge representation systems," in *Principles of Semantic Networks: Explorations in the representation of knowledge.*, J. Sowa, Ed. San Mateo, CA: Morgan Kaufmann, 1991, pp. 385-400.

[9] F. Baader, H.-J. Heinsohn, B. Hollunder, J. Muller, B. Nebel, W. Nutt, and H.-J. Profitlich, "Terminological Knowledge Representation: A Proposal for a Terminological Logic," DKFI, Technical Memo TM-90-04, 1991 1991.

[10] J. Sowa, *Conceptual Structures: Knowledge Representation in Mind and Machine*. New York: John Wiley & Sons, 1985.

[11] P. Zweigenbaum, B. Bachimont, J. Bouaud, and J. Charlet, "Issues in the structuring and acquisition of an ontology for medical language understanding," *Methods of Information in Medicine*, vol. 34, pp. 15-24, 1995.

[12] U. Hahn, M. Romacker, and S. Schulz, "How knowledge drives understa5nding - matching medical ontologies with the needs of medical language processing," *Artificial Intelligence in Medicine*, pp. (in press), 1998.

[13] R. Brachman and S. JB, "An overview of the KL-ONE knowledge representation system," *Cognitive Science*, vol. 9, pp. 171-216, 1985.

[14] R. J. Brachman, D. L. McGuinness, P. F. Patel-Schneider, L. A. Resnick, and A. Borgida, "Living with Classic: When and how to use a KL-ONE-like language," in *Principles of Semantic Networks: Explorations in the representation of knowledge*, J. Sowa, Ed. San Mateo, CA: Morgan Kaufmann, 1991, pp. 401-456.

[15] R. MacGregor, "Inside the LOOM description classifier," *SIGART Bulletin*, vol. 2, pp. 88-92, 1991.

[16] B. Nebel, "Computational Complexity of Terminological Reasoning in Back," *Artificial Intelligence*, vol. 34, pp. 371-383, 1988.

[17] N. Guarino and P. Giaretta, "Ontologies and knowledge bases: towards a terminological clarification," in *Towards Very Large Knowledge Bases*, N. Mars, Ed. Amsterdam: IOS Press, 1995, pp. 25-32.

[18] W. Nowlan, S. Kay, A. Rector, B. Horan, and A. Wilson, "PEN&PAD: A multi-lingual patient care workstation based on a unified representation of the medical record and medical terminology," Proceedings of Medical Informatics Europe (MIE-91), Vienna: Springer-Verlag, pp. 1043-1048, 1991.

[19] J. Kirby and A. L. Rector, "The PEN&PAD Data Entry System: From prototype to practical system," Proceedings of AMIA Fall Symposium, Washington DC: Hanley and Belfus, Inc, pp. 709-713, 1996.

[20] J. Rogers and A. L. Rector, "Terminological Systems: Bridging the generation gap," *Journal of the American Medical Informatics Association*, vol. (AMIA Fall Symposium Special Issue), pp. 610-614, 1997.

[21] J. Rogers, D. Solomon, A. Rector, and P. Zanstra, "From rubrics to dissections to GRAIL to classifications," Proceedings of Medical Informatics Europe (MIE-97), Thesalonika, Greece: IOS Press, pp. 241-245, 1997.

[22] E. Galeazzi, A. R. Mori, F. Consorti, A. Errera, and P. Merialdo, "A cooperative methodology to build conceptual models in medicine," Proceedings of Medical Informatics Europe '97, Porto Carras, Greece: IOS Pressq, pp. 280-284, 1997.

[23] E. Galeazzi, A. Rossi Mori, F. Consorti, and A. Errera, "MEDITA: the reference methodology for cooperative production of dissections by terminological analysis," GALEN-IN-USE Internal Report GIU/MET-1/97 version 1.1 (available from GALEN) 1997.

[24] A. Rossi Mori, A. Gangemi, G. Steve, F. Consorti, and E. Galeazzi, "An ontological analysis of surgical deeds," Proceedings of Artificial Intelligence in Medicine Europe (AIME-97), Grenoble, France, 1997.

[25] J. Cimino and P. D. Clayton, "Coping with changing controlled vocabularies," Proceedings of Eighteenth Annual Symposium on Computer Applications in Medical Care, Washington, DC: Hanley & Belfus, Inc, Philadelphia PA, pp. 135-139, 1994.

[26] K. Campbell, S. Cohn, C. Chute, G. Rennels, and E. Shortliffe, "Gálapagos: Computer-based support for evolution of a convergent medical terminology," Proceedings of AMIA Fall Symposium, Washington DC: Hanley and Belfus, Inc, pp. 269-273, 1996.

[27] D. E. Oliver, Y. Shahar, E. H. Shortliffe, and M. A. Musen, "Representing change in controlled medical vocabularies," *Artificial Intelligence in Medicine*, 1998.

[28] W. Solomon and H. Heathfield, "Conceptual modelling used to represent drug interactions," Proceedings of Twelfth International Congress of the European Federation for Medical Informatics, MIE-94, Lisbon, Portugal, pp. 186-190, 1994.

[29] A. Rector, S. Bechhofer, C. Goble, I. Horrocks, W. Nowlan, and W. Solomon, "The GRAIL concept modelling language for medical terminology," *Artificial Intelligence in Medicine*, vol. 9, pp. 139-171, 1997.

[30] J. Rogers and A. Rector, "The GALEN ontology," Proceedings of Medical Informatics Europe (MIE 96), Copenhagen: IOS Press, pp. 174-178, 1996.

[31] A. Rector, A. Gangemi, E. Galeazzi, A. Glowinski, and A. Rossi-Mori, "The GALEN CORE Model Schemata for Anatomy: Towards a re-usable application-independent model of medical concepts," Proceedings of Twelfth International Congress of the European Federation for Medical Informatics, MIE-94, Lisbon, Portugal, pp. 229-233, 1994.

[32] I. Horrocks and A. Rector, "Using a description logic with concept inclusions," Proceedings of International Description Logics Workshop (DL 96), Cambridge, MA: AAAI Press, Menlo Park, CA, pp. 132-135, 1996.

[33] I. Horrocks, "Comparison of Two Terminological Knowledge Representation Systems," : University of Manchester, 1995.

[34] I. Horrocks, A. Rector, and C. Goble, "A description logic based schema for the classification of medical data," Proceedings of Knowledge Representation Meets Data Bases (KRDB 96), Budapest, Hungary, pp. 24-28, 1996.

[35] I. Horrocks, "Optimising Tableaux Decision Procedures for Description Logics," in *Computer Science*. Manchester: University of Manchester, 1997, pp. 176.

[36] A. Rossi-Mori and CEN Project Team PT003, "Medical Informatics - Categorial structure of systems of concepts - Model of Representation of Semantics (ENV12264)," CEN TC251 WG2, ENV ENV 12264, 1995 1995.

[37] J. Bateman, B. Magnini, and G. Fabris, "The generalized upper model knowledge base: Organization and use," in *Towards Very Large Knowedge Bases*, N. Mars, Ed. Amsterdam: IOS Press, 1995, pp. 60-72.

[38] P. Zanstra, W. Claassen, and E. v. d. Haring, "A natural language generator in Rois," *LINKS*, pp. http://ehm.kun.nl/mi/project/rois/newsltr/Links2/Paper4/paper4.htm, 1997.

[39] J. Wagner, A.-M. Rassinoux, R. Baud, and J.-R. Scherrer, "Generating noun phrases from a medical knowledge representation," Proceedings of Twelfth International Congress of the European Federation for Medical Informatics, MIE-94, Lisbon, Portugal, pp. 218-223, 1994.

[40] W. A. Nowlan, "Clinical workstation: Identifying clinical requirements and understanding clinical information," *International Journal of Bio-Medical Computing*, vol. 34, pp. 85-94, 1994.

[41] R. H. Baud, J.-M. Rodrigues, J. C. Wagner, A.-M. Rassinoux, C. Lovis, P. Rush, B. Trombert-Paviot, and J.-R. Scherrer, "Validation of concept representation using natural language generation," *Journal of the American Medical Informatics Association*, vol. (Fall Symposium Supplement), pp. 841, 1997.

[42] Y. Shahar, A. K. Das, S. W. Tu, and M. A. Musen, "Knowledge-based temporal abstraction in clinical-management tasks," Proceedings of AAAi-94 Spring Symposium on Artificial Intelligence in Medicine, Stanford CA, 1994.

[43] F. Baader, "Augmenting concept languages by transitive closure of roles: An alternative to terminological cycles," Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-91): Morgan Kaufmann, pp. 446-451, 1991.

[44] B. R. Gaines, M. L. Shaw, and J. B. Woodward, "Modelling as Framework for Knowledge Acquisition Methodologies and Tools," *International Journal of Intelligent Systems*, vol. 8, pp. 155-168, 1993.

[45] World Health Organisation, *International Classification of Diseases*. Geneva: World Health Organisation, 1989.

[46] R. Côté and D. Rothwell, *SNOMED-International*. Chicago: College of American Pathologists, 1993.

[47] J. Read, "The Read Clinical Classification," NHS Centre for Coding and Classification, Loughborough, UK 1993.

[48] J. A. Bateman, "Upper modelling: a general organization of knowledge for natural language processing.," USC/Information Sciences Institute, Penman Development Note 1989.

[49] W. Ceusters, J. Rogers, F. Consorti, and A. Rossi-Mori, "Syntactic-semantic tagging as a mediator between linguistic representations and formal modals: an exercise in linking SNOMED to GALEN," *Artificial Intelligence in Medicine*, 1998.

[50] A. Rossi Mori, F. Consorti, E. Galeazzi, and P. Merialdo, "A second generation of terminological systems is coming," Proceedings of Medical Informatics Europe '97, Porto Carras, Greece: IOS Press, pp. 436-440, 1997.

[51] D. B. Lenat and R. V. Guha, *Building Large Knowledge-Based Systems: Representation and inference in the Cyc Project*. Reading, MA: Addison-Wesley, 1989.

[52] F. Giunchiglia and R. Sebastiani, "A SAT-based decision procedure for ALC," Proceedings of Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR'96), Boston, Mass: Morgan Kaufmann, pp. 304-314, 1996.

[53] W. Ceusters and P. Spyns, "From natural language to formal language: when MultiTALE meets GALEN," Proceedings of Medical Informatics Europe '97: IOS Press, pp. 396-400, 1997.

[54] C. Lovis, R. Baud, P.-A. Michel, and J.-R. Scherre, "Morphosemantic decomposition and semantic representation to allow fast and efficient natural language recognition," *Journal of the American Medical Informatics Association*, vol. (Fall Symposium Supplement), pp. 873, 1997.

[55] A.-M. Rassinoux, "Modeling just the important and relevant concepts in medical language understanding," *Methods of Information in Medicine*, pp. (in press), 1997.

[56] R. Baud, C. Lovis, L. Alpay, A.-M. Rassinoux, J.-R. Scherrer, A. Nowlan, and A. Rector, "Modelling for Natural Language Understanding," Proceedings of Seventeenth Annual Symposium on Computer Applications in Medical Care (SCAMC-93), Washington DC: McGraw Hill, pp. 289-93, 1993.

[57] H. Gu, Y. Perl, and J. Geller, "A methodology for partitioning a vocabulary hierarchy into trees," *Artificial intelligence in Medicine*, pp. (in press), 1998.

[58] L. Padgham and P. Lambrix, "A Framework for Part-of Hierarchies in Terminological Logics," Proceedings of KR-94, pp. 485-96, 1994.

[59] J. Bernauer and H. Goldberg, "Compositional classification based on conceptual graphs," Proceedings of Artificial Intelligence in Medicine Europe (AIME-93), Munich, 1993.

[60] J. Smart and M. Roux, "Medical knowledge representation for medical report analysis," Proceedings of Fifth conference on Artificial Intelligence in Medicine Europe (AIME '95), Pavia, Italy: Springer, pp. 53-66, 1995.

[61] E. Franconi, G. DeGiacomo, R. M. MacGregor, W. Nutt, C. A. Welty, and F. Sebastiani, "International Workshop on Descripion Logics (DL '98)," Proceedings , Trento, Italy: http://anemosyne.itc.it/1024/dl/dl98, 1998.

Acknowledgements

This work supported in part by grants A2012 and HC2018 from the Healthcare Telematics Programme of the European Commission. Figure 1 courtesy of Semantic Technologies Limited, Enterprise House, Manchester Science Park, Manchester M13 6SE. Details of all members of the GALEN and GALEN-IN-USE consortia may be found at http://www.cs.man.ac.uk/mig/galen .

Appendix: Brief informal description of relevant constructs in the GRAIL representation language.

The following describes the relevant subset of GRAIL syntax and semantics in terms of first order logic, since standard description logic notation may be unfamiliar to many readers, and those familiar with it should be able to make the appropriate translations.

A GRAIL concept model consists of a set of categories ('concepts') *C* and a set of Attributes *A*. There is a subset of categories EC $\subseteq$ C. The semantics are defined by interpreting categories as predicates over a model of individuals. The basic operators relevant to this paper are given below.

C_i <u>*newSub*</u> EC_j	Introduce a new elementary category EC_j subsumed by E_i
EC_i <u>*addSub*</u> EC_j	Elementary category EC_i subsumes Elementary category EC_2.
C <u>*which*</u> $<A_1C_1\ A_2C_2 \ldots A_nC_n>$ (<…> may be omitted if *n=1*)	Introduce a new anonymous concept C_a such that $"\ x\ .\ C_a x\ «\ \ Cx\ \&\ (\$y_1\ .\ \ C_1y_1\ \&\ xA_1y_1)\ \&\ (\$y_2\ .\ \ C_2y_2\ \&\ xA_2y_2)\ \&\ \ldots\ \&\ \ (\$y_n\ .\ \ C_ny_n\ \&\ xA_ny_n)$
String = C (or *C* <u>*name*</u> *String)*	The string *String* is an external name, or abbreviation, for the concept *C*.
A_1 <u>*newAttribute*</u> $A_2\ A_3$ *cardinality*	Introduce a new attribute A_2 subsumed by A_1, with inverse A_3 and with *cardinality* $\in$ *{ oneOne manyOne oneMany manyMany}*
A <u>*transitiveDown*</u>	Attribute A is transitive. (The 'down' is a useful hint to the software for optimisation.)
A_1 <u>*specialisedBy*</u> A_2	$"\ \ x\ y\ z.\ xA_1y\ \&\ yA_2z \rightarrow xA_1y$ A_2 *must be transitive.*