Ontological Issues in using a Description Logic to Represent Medical Concepts:
Part II - The GALEN High Level Schemas

A.L Rector

J.E. Rogers

Medical Informatics Group
Department of Computer Science
University of Manchester
Manchester M13 9PL
Tel +44-161-275-6188/6239/7183
FAX: +44-161-275-6204
email <rector/jeremy>@cs.man.ac.uk

& *Open*GALEN
www.opengalen.org

Abstract

A previous paper [1] described the principles behind the development of the GALEN ontology and its high level schema. This paper describes the high level schemas themselves for the top level, anatomy and disease, and illustrates how those principles have been worked out in practice. The complete Common Reference Model and associated material is available from *Open*GALEN at www.opengalen.org

Keywords: Terminology, Ontology, Knowledge Representation, conceptual modelling

1. Introduction

GALEN seeks to provide re-usable terminology resources for clinical systems. The heart of GALEN is an ontology, the Common Reference Model, formulated in a specialised description logic, GRAIL [2]. The key choices behind the ontology are embodied in its high level schemas – the choice of major categories, semantic links, and representation style for key constructs. A previous paper [1] describes the rationale behind GALEN's high level schemas. This paper describes the top level schema plus the high level schemas for anatomy and disease of the GALEN internal representation. Separate papers describe aspects of the schemas for procedures [3] and drugs [4].

Although presented in GALEN's notation, the presentation is intended to be sufficiently general to allow comparison and potential harmonisation with other clinical ontologies such as that being developed by SNOMED-RT [5, 6], the UK Clinical Terms project [7], the Digital Anatomist Project [8], or more language oriented work such as that of Zweigenbaum [9] or Hahn [10] or the ontologies used for problem solving by Musen [11]. The complete ontology along with documentation can be found on the *Open*GALEN web site www.opengalen.org.

It must be emphasised that these are the schemas for the underlying ontology formulated in the GRAIL description logic. GALEN treats this ontology as an 'assembly language' which few users ever see. The goal of this underlying ontology is to be unambiguous to result in correct classification. Intuitive presentations to users are dealt with separately [12-14].

2. Top Level Ontology

GALEN's top level distinction is between first class entities, or *Things*, and everything else, or *ModifierConcepts*. *Things* are divided into

- *GeneralisedStructures* — abstract or physical things with parts independent of time
- *GeneralisedSubstances* — continuous abstract or physical things independent of time
- *GeneralisedProcesses* — changes which occur over time

This structure is adapted from Lenat and Guha. However Lenat and Guha maintain a distinction for processes analogous to that between 'Structure' and 'Substance', *e.g.* between "The digestion of a meal" and "The activity of digestion". In GALEN's experience, both users and knowledge engineers have found this distinction confusing, and it was dropped.

Too many medical processes have ill-defined beginnings and ends. Similarly, the naming of the category 'Thing' led to arguments so it was left implicit. The GALEN Common Reference Model as published is therefore as shown in Figure 1.

2.1 Modifiers

The first level break down of *ModifierConcept* falls into:

- Aspect and Modality
  - *Aspect* –'modifiers' proper, normally used in the pattern: *PrimaryThing which …modifier* (see [1])
  - *Modality* – modalities as described under *First class entities and modifiers* normally used in the pattern *Modality which …* (see [1])
  - Other concepts which are dependent on first class concepts for their meaning
  - *Role* – arbitrary concepts used to make elementary taxonomies orthogonal, *e.g. DoctorRole, HormoneRole, DrugRole* etc. See *Roles and role designating attributes,* below.
  - *Collection* – set, system, etc. GALEN has a weak notion of *Collection*. There are no special features in the formalism to support operations on collections.
- Other things that have special significance or behaviour
  - *Unit – mg, day, etc.*

Aspect is further broken down into:

- *Features* – nominalised relations between such as 'Level' or 'Severity' which must be modified by one or more States to have a meaning in a 'Feature-State' pair
- *States* – the 'values' of *Features, e.g.* 'mild', 'moderate', or 'severe'
- *Selectors* – values of selections, largely in anatomy, such as 'left/right' or 'upper/middle/lower' as opposed to 'to the left of' and 'to the right of', etc. Selectors identify a specific entity rather than modifying it.
- *Statuses* – modifiers used in the 'internal' workings of the model such as *'normal/nonNormal', 'countable/indefinitelyDivisible/mass'* and various topological indicators.

All mutable properties are nominalised to *feature-state* pairs, *e.g. Disease which hasFeature (Severity which hasState severe).* By contrast, *Selectors*, and *Statuses* are always linked to the entity they modify by a single semantic link, *e.g. Hand* <u>*which*</u> *hasLeftRightSelector rightSelection.* There can never be anything more to say about the 'rightness' of the right hand, it just is.

2.2 Phenomenon – Secondary High Level Ontology

As noted in [1] most ontologies seem to require at least one very high level disjunctive category to represent key concepts. In GALEN this is *Phenomenon.* A *Phenomenon* is anything which can be, or can be modified to be, worth noting clinically as nonNormal or pathological. As shown in Figure 2, *Phenomenon* is the disjunction of *Thing (GeneralisedStructre, GeneralisedProcess, GeneralisedSubstance), Feature, Status,* and *Collection*. This is really too inclusive for a general ontology, since it allows things like rocks to be 'pathological', but the enormous effort to tighten the constraints has not so far been warranted in a clinical ontology for clinical applications.

3. High level attribute ontology

3.1 Top Level attribute ontology

The top level breaks down into:

ConstructiveAttribute -	Links between first class entities, *i.e. Things, i.e. GeneralisedStructures, GeneralisedSubstances,* and *GeneralisedSubstances*
ModifierAttribute -	Links between *Things* and *ModifierConcepts*
TemporalAttribute -	Links between processes and statuses via temporal relations (not fully complete)

3.2 Constructive Attributes

The key constructions in most medical terms are to locate a disease or procedure in a structure or part of a structure. The most common schema in GALEN for disease or procedure concepts, and by far the most common schema for queries or abstractions is:

Disease/Procedure <u>*which*</u> *LocativeAttribute BodyStructureOrProcess*

LocativeAttribute has been steadily generalised in the course of the project until it has become the analogue of 'Phenomenon', a domain specific disjunction of attributes needed for high level queries.

ConstructiveAttributes therefore break down into three main attributes plus the domain specific disjunction, *LocativeAttribute (*alias *involves)*:

PartitiveAttribute - Part-whole relations as discussed further under anatomy and processes.

StructuralAttribute - Non-partitive structural relations such as connects, passes through, contains, etc.

FunctionalAttribute - Attributes involving functional relations.

involves (LocativeAttribute) - Links lesions, processes, and procedures to their location or cause. Tumours are 'located in structures' rather than being 'part of structures' in the GALEN ontology.

It may seem odd to think of causal attributes as 'locative', but, for example, the classification of 'spider haemangiomata' under 'Phenomenon involving liver disorder' is appropriate. .

3.3 Modifier Attributes

The modifier attributes and modifier categories are intimately tied, one main branch of the attribute hierarchy for each branch of the modifier hierarchy: *hasFeature, hasState, hasStatus, hasSelector, hasModality*.

3.4 Roles *and* RoleDesignatingAttributes

The principle of orthogonal taxonomies leads to a wide range of specialised hierarchies which are used to designate aspects of concepts which themselves are best seen as natural kinds. In GALEN these are all known as 'roles' and linked by attributes named playsXRole or hasXRole. The range of usage extended from marking substances as hormones using *homoneRole* to identifying professions by playsClinicalRole.

4. Anatomy

4.1 Part-whole relations and physical connection

There has been much study of parts and wholes, or 'partitive relations', and an entire field of 'merology' is based on their study [15, 16] and significant work in the description logic community [17-19].

GALEN's principles are:

a) All partitive relations are transitive.

b) Diseases/disorders/procedures of/on a part pertain also to the whole.

c) Partitive relations break down into two groups: those involving parts of physical objects and those involving constituents in mixtures of substances or granular material. (Collections are not considered ‘partitive’ in GALEN though they are in [15, 16])

d) Partitive relations involving physical parts occur in three main ways:

- Components – things like joints, ligaments, processes, organs, which occur only in one (or occasionally more) divisions of an object. Components of any sort of part are components of the whole.
- Solid and Surface divisions – things which are roughly self similar, at least to the extent that they have similar layers, *e.g.* the upper and lower arm are divisions of the upper extremity. The distinction between solid and surface divisions is roughly parallel to the distinction between two- and three-dimensional regions in the Digital Anatomist Project [8]. Divisions of components are not components of the whole.
- Layers – things like skin or the muscle or periosteum that occur in all divisions of an object. GALEN’s approximation is that a layer of a division is a kind of a layer of the whole. The more correct relation as pointed out in [17] is that a layer of a division should be a division of a layer of the whole – *i.e.* the skin of the hand should be a division of the skin of the upper extremity. Layers of components are not layers of the whole – *i.e.* the surface of the cusp of a heart valve is neither a kind of nor a part of the surface of the heart.

e) Connection is transitive but not partitive; ‘branch-of’ is neither partitive nor transitive (otherwise all arteries would be branches of the aorta – perhaps true in some sense but not very useful)

f) Connected physical sets such as the ‘digestive tract’ are distinct from functional systems such as ‘the digestive system’. For example, the pancreas is part of the digestive system but not of the digestive tract.

4.2 Topology, cavities and containment

All structures in the GALEN Common Reference Model have a topology that may be *hollow* or *topologicallySolid*. Being solid is simple. The key feature of being hollow is that any hollow object *definesSpace* known as a *Cavity*. Things can only be ‘contained’ in the cavities formed by hollow objects. Containment is not considered universally partitive in the GALEN Common Reference Model because it is not consistently the case that ‘disorders of a contained thing are a disorder of the containing thing' – "disorders of the heart" are not usually considered "disorders of the pericardium' (in fact, if anything, the opposite is true).

Being *hollow* is actually quite complex and breaks down into

surfaceHollow – Surface regions such as the 'abdomen' which overlie a cavity and are often said to have things in them. This may be an artificial convenience but it is difficult to cope with many common medical expressions without some such construct.

trulyHollow – Properly hollow structures.

closedHollow – No openings

tubularHollow – One or two openings. The cavity is a *lumen*.

bilayered – Membranes which form potential spaces such as the peritoneum or pleura

4.3 *Tissues, Cells, and substances:* mass, discrete, *and* indefinitelyDivisible

Most western languages make a distinction between 'mass nouns' such as which are normally used in the singular – such as "water", "sand", and "shopping" – and 'count nouns' that may be singular or plural– such as "stick", "stone", and "purchase". Lenat and Guha make a corresponding semantic distinction *e.g.* between 'stuff' and 'thing' [20].

In the GALEN Common Reference Model, structures and substances have a *Countability* that can be one of:

discrete – *individual bones, organs, membranes, etc.*

mass – *substances and tissues*

indefinitelyDivisible– *Cells, grains of sand, etc.*

The *indefinitelyDivisible* category was added to cover things like cells which are usually treated en masse as in their count-concentration in a body fluid, but which can have discrete parts.

The partitive attribute for mass and indefinitelyDivisible attributes are *isMadeOf*. The containment (non-transitive) attribute is *isMixedThroughout* .

4.4 *Regions*

The problem of describing the regions of the body is one of the significant headaches for any system attempting to describe anatomy logically.

a) Regions have ill defined borders

b) Regions can be either two-dimensional surface regions or three-dimensional solid regions, and the distinction is not always clear.

c) Regions are often named identically with the primary structure which they contain: the region of the lower extremity associated with the “knee joint” is the “knee region”.

d) The thing for which a region is named may be based on clinical significance and cannot be inferred, *e.g.* there are many structures in the left anterior chest, but the “precordium” is specifically associated with the heart.

e) The definitions of some regions, such as the axilla and perineum, varies amongst authorities.

GALEN represents two- and three-dimensional divisions of regions using two families of attributes – *hasSolidDivision* and *hasSurfaceDivision*. (*i.e.* solid/3-dimensional regions) and correspondingly *SolidRegions* and *SurfaceRegions*. Regions are typically named for a single specific body part, such that the “periaortic space” is described by *hasSpecificProximity Aorta.*

4.5 Bits and Pieces

Terms such as ‘capsule’, ‘spine’ or ‘edge’ are widely used in anatomy to identify anatomical substructure elements –*e.g.* “capsule of kidney”, “spine of 5th lumbar vertebra”, “edge of liver” etc. Each such term can be ascribed at least some level of definition although such definition may be less than precise.

In modelling such anatomical substructure there are two choices.

a) To represent the generic notions as natural kinds and the real anatomic structures as defined composites, *e.g. Angle* <u>*which*</u> *isStructuralComponentOf Mandible, Pole* <u>*which*</u> *isStructuralComponentOf Kidney, etc.*

b) To represent the substructures that can themselves be taken as natural kinds, *e.g. PoleOfKidney, PoleOfOvary*, etc. with no explicit relationship to the more abstract notion of e.g. *Pole*

In general, GALEN has chosen the compositional method because there seems to be sufficient commonality in notions such as “lobe”, “pole”, “segment”, etc. to merit capturing them individually.

4.6 Other anatomical notions represented

In addition to whether objects are solid or hollow, mass or discrete, GALEN represents three other anatomical notions: *SurfaceVisibility* – whether a structure is internal or external; *Shape* – laminar, linear, etc.; and *pairing* – paired or unpaired.

4.7 Breaking up long lists and the NAMED… *convention*

The principle of orthogonal taxonomies, combined with a principled approach to choosing natural kinds, results in a broad flat hierarchy of elementary categories which is difficult to work with. For convenience GALEN breaks this hierarchy up into units which are convenient for the developers by introducing categories such as *NAMEDBodyPart*, *NAMEDDrug*, etc.

5. Diseases

5.1 Normal/NonNormal and Physiological/Pathological

What is a "disease" or "disorder"? What does it mean to say that something is "normal" or "abnormal"? "pathological" or "physiological"? Given many different philosophical definitions, the only recourse was to identify what operational outcomes were required to provide a reasonable logical approximation. These include:

a) Distinguishing normal anatomy from abnormal and to list the normal anatomical parts, connections, etc. of any structure.

b) Identifying those things whose presence was potentially noteworthy in a medical record or similar.

c) Flagging things as clearly 'diseases' or 'pathological', *i.e.* something close to "in potential need of medical management"

d) Representing the notion of being 'abnormal but not pathological' which we took as meaning something like "note-worthy but not in need of medical management".

e) Recognising that the presence of some things is always pathological, *e.g.* a malignant tumour or a fracture

To achieve these objectives, GALEN defines two independent status distinctions, *normal/nonNormal* and *pathological/physiological* plus two rules enforced through GRAIL's necessary statement mechanism:

a) anything *pathological* is *nonNormal*, and

b) anything *normal* is *physiological*

In addition it defines a further set of statuses: *intrinsicallyNormal/intrinsicallyNonNormal* and *intrisicallyPhysiological/intrinsicallyPathological* plus the rules:

a) Anything *intrinsicallyNonNormal* is *nonNormal*

b) Anything *intrinsicallyPathological* is *pathological*

The converses are not true, things which are *intrinsicallyNormal* are not necessarily *normal* – they may have acquired disorders or abnormalities; likewise for *intrisicallyPhysiological* and *physiological*.

5.2 Phenomenon *and "Disease"*

What then is a "disease"? The closest logical approximation to "Disease" or "Disorder" in the GALEN Common Reference Model is *PathologicalPhenomenon* defined as:

Phenomenon <u>which</u> hasPathologicalStatus pathological

The closest logical equivalent to "Disease of Organ or System" is:

PathologicalPhenomenon <u>which</u> involves OrganOrSystem

For example, "cardiovascular disease" is represented by the GALEN concept *CardiovascularPathology* which is defined as:

PathologicalPhenomenon <u>which</u> involves CardiovascularSystem

The GALEN category *Phenomenon* and attribute *involves* have been carefully crafted to try to capture the various ways in which things can go wrong with organs or processes to constitute "diseases of …" or "disorders of …". The label *PathologicalPhenomenon* explicitly avoids implying too close a mapping to any natural language phrase such as "disease", "disorder", or "condition".

5.3 Causation

Causation is a critical notion to medical concepts but surprisingly slippery. The GALEN Common Reference Model recognises roughly four dimensions.

a) Strength of association – from statistical to physiological cause

b) Immediate vs late

c) Whether thought of as a 'complication' or a 'cause'

d) Whether conceptualised as the primary or secondary cause as indicated by whether it is used in the naming or not.

Attributes which indicate close causal connections, *e.g.* *isSpecificImmediateConsequenceOf* are transitive whereas attributes indicating loose connections such as *isLateConsequenceOf* or *hasAssociation* are not.

5.4 Multiple causation

Many conditions are defined by their cause, *e.g.* "viral pneumonia", "bacterial meningitis", etc. What is to be done about situations in which there is more than one cause? Clinicians dislike the logical inference that "mixed pneumonia" is a kind of "bacterial pneumonia" and at the inability to distinguish between a "mixed pneumonia" and a "viral pneumonia complicated by bacterial infection". Nor would any clinician want to imply that "diabetic renal failure" could not have other contributing causes. On the other hand, the notion of an "infection involving bacteria" certainly ought to include such mixed conditions.

GALEN's solution to this problem is to use special child attributes of *isImmediateConsequenceOf*, *isLateConsequenceOf* *i.e.* *isSpecificImmediateConsequenceOf*, *isSpecificLateConsequenceOf*, etc. , *e.g.* *ViralPneuomnia* is defined as *Pneumonia* <u>*which*</u> *isSpecificImmediateConsequenceOf ViralInfection.*

5.5 Levels

A recurrent pattern pointed out by Shahar [21] is in handling departures from the expected or normal value and changes in states. GALEN has adapted Shahar's scheme to provide a consistent ontology for all measurements which can be elevated, depressed etc. based on the re-usable quantity *Level* which takes a series of attributes linking it to different kinds of state: *hasMagnitude, hasChangeInState, hasTrendInState, hasRelativeLevelState, and hasExpectedLevelState,* allowing complex ideas such as a "temperature with a magnitude of 39C, which is falling but considered elevated (higher than expected)" .

5.6 Clinical Situations, presence and absence – encapsulating concepts for Medical Records

Two of GALEN's specific objectives as indicated in the introduction were to encapsulate concepts so that they could be incorporated into traditional medical record architectures and to provide means of mapping to existing coding and classification schemes. What is it that must be entered in a record or that must be mapped to a coding scheme? Answers to the two questions are similar but not identical.

For medical records which have been designed for traditional coding schemes:

a) It must be able to include the negative as well as positive concepts, since some medical records systems provide negation within their information model but others do not.

b) It is often a complex of several conditions – *e.g.* A with B without C

c) It is often necessary to record causal or temporal relations to other things in the medical record and co-ordinate that information with the concept – *e.g.* “nephropathy secondary to diabetes” or “fat embolism secondary to fracture of femur”.

GALEN achieves a) and b) by ‘wrapping’ the kernel concept in two outer modalities:

a) *Existentiality – presence* or *absence*

b) *ClinicalSituation* – a clinical ‘chunk’ to be recorded and treated together.

Hence the total expression for a concept representation for recording in a medical record or mapping to a coding system is always of the form of the example below:

ClinicalSituation which isCharacterisedBy <
presence which isExistenceOf StomachUlcer
presence which isExistenceOf StomachPenetration
absence which isExistenceOf Haemorrhage>.

Transitivity is declared in such a way that a *ClinicalSituation* which is characterised by the *presence* or *absence* of some condition, is itself characterised by that condition – hence the above example would be subsumed by both *ClinicalSituation which involves Stomach* and *ClinicalSituation which involves Haemorrhage.*

5.7 Mapping to coding and classification systems

In general, ‘clinical situations’ map directly to items in traditional coding and classification systems such as ICD9/10 or the disease axis of SNOMED-International with a few provisos.

a) The categories in the GALEN Common Reference Model do not represent codes directly. Rather they are mapped to codes using the indexing methods described [22]. The general rule is that a GALEN concept is mapped to the most specific code on which it is indexed. If there is more than one such code, then other mechanisms for disambiguation are required. In general “excluding” clauses in ICD –. *e.g.* “hypertension excluding pregnancy” – indicate that a more specific code exists elsewhere and are catered for automatically by this mechanism.

b) A code may be mapped to more than one GALEN category, typically if there is an “includes” clause in its rubric.

c) Any code rubric including ‘other’ or ‘not elsewhere classified’ is mapped to an artefactual GALEN concept specially created for that coding system.

d) Any code rubric including ‘Not otherwise specified’ (‘NOS’) is mapped to the parent concept with a suitable annotation on the mapping.

6. Discussion

The GALEN ontology schemas illustrate how the principles set out in the previous paper are carried through in practice in the internal form of the GALEN ontology. The fundamental criteria are a) the ability to express medical concepts and the abstractions required for medical queries and b) the correct classification and retrieval of those concepts subject to the known limitations identified [1]. Human factors issues are dealt with separately in GALEN’s mechanisms to support user interfaces [12].

With respect to the internal criteria previously identified [1], the representation is sufficiently expressive that few instances arise where the authors must ‘lie to the system’, although the model is still clearly at best a ‘logical approximation’, and the labels, particularly on semantic links, give only an approximate indication of their use. There are both distinctions that have no linguistic equivalence and also linguistic subtleties that are not catered for. The use of *presence/absence* and similar constructs to convey negation is conservative, but limits expressiveness only rarely in the applications we support. The structure of transitivity avoids ‘up-to-ism’ and allows concepts to be represented consistently regardless of context (‘concept constancy’). The major issues of pragmatic normalisation occur in the procedure and drug schemas not covered in this paper [4].

With regards to external criteria for expressiveness and classification, there are two questions:

a) Does the ontology meet the criteria? *i.e.* Is the ontology sufficiently expressive? Does it result in correct classification?

b) Is the ontology over-engineered? *i.e.* could the same result be obtained more simply?

Evidence for a) comes primarily from the surgical and drug extensions to the ontology where the ontology has so far proved adequately expressive to formulate the concepts in existing terminologies and classifications. Preliminary comparisons with other classifications have been undertaken for procedures [23] and further evaluations are an ongoing part of the development of the development of the drug ontology [4]. Informal comparisons with the Digital Anatomist project [8] (www1.biostr.washington.edu/DigitalAnatomist.html) suggest a high level of concurrence in the principles of the respective anatomic models. However much wider discussion is required which is best achieved by comparison with other classifications.

As to the question of whether GALEN's ontology is over-engineered, this can only be answered by demonstrating that similar results can be obtained more simply, which similarly requires comparison with other ontologies.

One of the major purposes of this paper and of making the ontology open source through *Open*GALEN is to facilitate such comparisons and development. The full ontology along with associated material is available at www.opengalen.org.

Acknowledgements

The research reported was supported in part by the European Healthcare Telematics programme. The partners in the GALEN and GALEN-IN-USE projects have contributed extensively. Particular mention is due to Anthony Nowlan who first formulated key parts of what has become the GALEN ontology.

References

1. Rector A, Rogers J. Ontological issues in using a description logic to represent medical concepts: Experience from GALEN. Methods of Information in Medicine 2000((in press)).
2. Rector A, Bechhofer S, Goble C, Horrocks I, Nowlan W, Solomon W. The GRAIL concept modelling language for medical terminology. Artificial Intelligence in Medicine 1997;9:139-171.
3. Rogers J, Solomon D, Rector A, Zanstra P. From rubrics to dissections to GRAIL to classifications. In: Medical Informatics Europe (MIE-97); 1997; Thesalonika, Greece: IOS Press; 1997. p. 241-245.
4. Solomon DS, Wroe C, Rogers JE, Rector A. A reference terminology for drugs. Journal of the American Medical Informatics Association 1999(1999 Special Conference Issue).
5. Spackman KA, Campbell KE, Côté RA. SNOMED-RT: A reference Terminology for Health Care. Journal of the American Medical Informatics Association (JAMIA) 1997((Symposium special issue)):640-644.
6. Spackman KA, Campbell KE. Compositional concept representation using SNOMED: Towards further convergence of clinical terminologies. Journal of the American Medical Informatics Association 1998(1998 Fall Symposium special issue):740-744.
7. Price C, Brown PJB, Schulz EB. Comparing Primitive Concept Fields in Clinical Terminologies. Journal of the American Medical Informatics Association 1997(1997 Fall Symposium Special Issue).
8. Rosse C, Shapiro IG, Brinkley JF. The Digital Anatomist foundational model: Principles for defining and structuring its concept domain. Journal of the American Medical Informatics Association 1998(1998 Fall Symposium Special issue).
9. Zweigenbaum P, Bachimont B, Bouaud J, Charlet J. Issues in the structuring and acquisition of an ontology for medical language understanding. Methods of Information in Medicine 1995;34(1/2):15-24.

10. Hahn U, Romacker M, Schulz S. How knowledge drives understanding - matching medical ontologies with the needs of medical language processing. Artificial Intelligence in Medicine 1999;15(1):25-52.

11. Eriksson H, Puerta A, Musen M. Generation of knowledge-acquisition tools from domain ontologies. International Journal of Human-Computer Studies 1994;41:399-424.

12. Rector AL, Zanstra PE, Solomon WD, Rogers JE, Baud R, Ceusters W, et al. Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine. IEEE Transactions on Information Technology in BioMedicine 1999;2(4):229-242.

13. Rector A, Rossi Mori A, Consorti F, Zanstra P. Making sound re-usable terminology practical: The GALEN approach. In: Waegemann P, editor. Towards an Electronic Health Care Record, Europe; 1998; London: Medical Records Institute, Newton Mass; 1998. p. 298-305.

14. Rector A. Thesauri and formal classifications: Terminologies for people and machines. Methods of Information in Medicine 1998;37(4-5):501-509.

15. Odell JJ. Six different kinds of composition. Journal of Object Oriented Programming 1994;5(8):10-15.

16. Winston M, Chaffin R, Hermann D. A taxonomy of part-whole relations. Cognitive Science 1987;11:417-444.

17. Padgham L, Lambrix P. A Framework for Part-of Hierarchies in Terminological Logics. In: Sandewall E, Torasso P, editors. KR-94; 1994; 1994. p. 485-96.

18. Artale A, Franconi E, Guarino N. Open problems for part-whole relations. In: International Workshop on Description Logics; 1996; Boston, MA; 1996. p. http://www.cs.man.ac.uk/franconi-bin/retrieve/links/franconi-papers/part-papers/part-dl96.dvi.gz.

19. Artale A, Franconi E, Pazzi L. Part-whole relations in object-centered systems: An overview. Data and Knowledge Engineering 1996;20:347-383.

20. Lenat DB, Guha RV. Building Large Knowledge-Based Systems: Representation and inference in the Cyc Project. Reading, MA: Addison-Wesley; 1989.

21. Shahar Y, Tu S, Das A, MA M. A problem-solving architecture for managing temporal data and their abstractions. In: Workshop on Implementing Temporal Reasoning, AAAI-92; 1992; San Jose, CA; 1992. p. 141-52.

22. Rector AL, Kay S. Foundations of the electronic medical record revisited. Methods of Information in Medicine 2000:(in preparation).

23. Rogers JE, Price C, Rector AL, Solomon WD, Smejko N. Validating clinical terminology structures: Integration and cross-validation of Read Thesaurus and GALEN. Journal of the American Medical Informatics Association 1998(Fall Symposium Special Issue):845-849.

Figure 1: Top Level GALEN Category Ontology

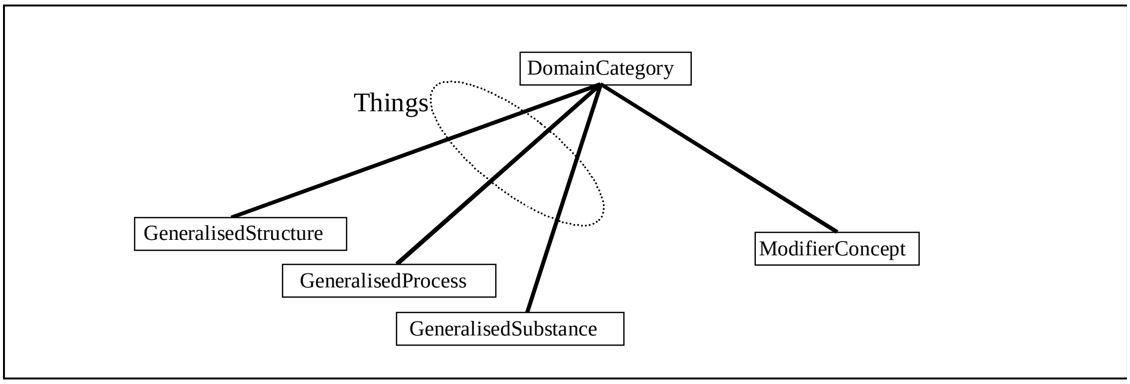

Figure 2 : Secondary High Level Taxonomy