Untangling Taxonomies and Relationships: Personal and Practical Problems in Loosely Coupled Development of Large Ontologies

Alan L. Rector, Chris Wroe, Jeremy Rogers, Angus Roberts

Medical Informatics Group
Department of Computer Science, University of Manchester
Manchester M13 9PL, UK
Tel: +44-161-275-6188/6239/7183 FAX:+44-161-275-6204
email {rector|wroec|jrogers|aroberts}@cs.man.ac.uk

Abstract

The GALEN prALogramme has been developing medical ontologies collaboratively for nearly a decade. The ontologies are large and formulated in a specialised description logic, GRAIL. The programme is a broad collaboration of over a dozen groups, most with no prior experience of developing formal ontologies. The programme has developed a methodology for loosely coupled development using layers of intermediate representations, guidelines and tools which minimises training requirements for domain experts and effort by central knowledge engineers.

Issues arise both from problems in formal representations and from the idiosyncrasies of the medical domain. Issues dealt with include 'tangled' taxonomies, part-whole and locative relationships, defaults and exceptions, semantic normalisation, and the difference between medical convention and strict logical criteria for correctness.

Keywords:

Cooperative development; ontology development; ontology design; very large ontologies, medical

INTRODUCTION

The GALEN programme has been developing medical ontologies collaboratively for nearly a decade. (The current versions are available through *Open*GALEN at www.opengalen.org.) The ontologies are large – over 20,000 surgical procedures, nearly 10,000 anatomical concepts and 3000 drugs – the schemas are complex – between fifty and one hundred families of link types covering different flavours of partonomy and location, function, and causation. They are formulated in a specialised description logic, GRAIL[15-17].

The programme is a loose collaboration which has varied over time from seven to more than a dozen groups. Most of those contributing to the development have little prior experience of building formal ontologies, although many have long experience of developing and or testing medical terminologies. Some were actively sceptical or even hostile to the ideas of formality and standardisation. In the testing phase of the programme, most groups were part of other larger efforts with a their own primary goals, so that the effort available for this project was limited. Training time for most workers had to be confined to no more than six days divided into two workshops.

The same methods have been applied to the development of a large ontology of drugs, their uses, actions, side effects, etc. as part of the UK PRODIGY and Drug Ontology projects [20, 31].

This paper describes the interplay of the different methodological and technical elements which have been brought to bear on this problem and the overall approach and rationale for ontology development which has emerged from it. We view the problems of ontology development as an intimate mix of organisational and technical issues in which different interests and priorities must be reconciled to achieve a successful outcome.

BASIC ELEMENTS OF ONTOLOGY CONSTRUCTION IN GALEN

Goals And Criteria for Correctness

*Open*GALEN aims to produce clinical ontologies which are:

- *Logically correct* and therefore suitable for use in retrieval, rule based systems, etc. For example, all and only "heart diseases" should be classified under "heart disease", all and only procedures on the liver under "liver procedures" etc. Any given concept should be classified in as many ways as appropriate.
- *Reusable* and therefore suitable support system integration, communication etc. The resulting classifications must therefore to contain as fine a grained

K-CAP'01, October 22-23,2001, Y Gil, M Musen, J Shavlik (eds) Victoria, British Columbia, Canada:139-146. Copyright 2001 ACM 1-58113-380-4/01/0010…$5.

detail and support as many alternative views as are required by the union of the applications that might reasonably be expected to use them.

By contrast most existing medical terminologies, with the exception of SNOMED-RT [22], have been designed for single applications – *e.g.* bibliographic retrieval, remuneration, or epidemiological reporting – and organised to facilitate access intuitive access by clinicians rather than logical correctness or accuracy of retrieval.

Basis of the approach

*Open*GALEN's requirements are for distributed loosely coupled development of complex ontologies with only modest need for central coordination and limited possibility of central control. The vast majority of the participants are interested in the outcomes rather than the underlying process.

There are four groups who need to participate in the development of an ontology:

- *Content* contributors
- *Domain experts* who capture and quality assure that bulk of the content formally, usually but not always based on some external source of content
- *Knowledge engineers* who design and maintain the formal ontology itself
- *Logicians* who develop and maintain the underlying logic engines and representations

A comprehensive methodology must coordinate the activities of all four groups and provide clean interfaces between them. However, *Open*GALEN concentrates on reconciling viewpoints of different domain experts – often distributed amongst many centres with many different priorities – with those of the knowledge engineers. This distributed approach has led us to a different emphasis from that of authors such as Uschold [26].

The goal has been to allow domain expert's to work as independently as possible with guidelines and agreements which are intuitive at a domain level while, at the same time, allowing the knowledge engineers maximum freedom to develop the underlying description logic ontology.

Domain experts therefore work in tailored 'intermediate representations' which are transformed algorithmically into the underlying description logic based ontology. The use of intermediate representations has a long history in knowledge based systems generally [3], and the use of schemas to create specialist environments for domain experts has many analogies with the approaches of PROTÉGÉ [11, 24, 25]and KADS[28]. However, it has been less widely used in ontology development, although Staab [23] describes the use of a somewhat lower level intermediate representation to separate developers from the details of implementation. Staab's intermediate representation is closer to the level of GRAIL or the rapidly developing interchange language, OIL[1] It addresses the issues of translating these relatively high level description logic into low level expressive representations such as FaCT [7, 8]. FaCT is proving difficult even for experienced knowledge engineers to use directly. However, our domain experts find even knowledge-engineering oriented languages at the level of GRAIL or OIL difficult to manage. We therefore envisage the continuing need both for a knowledge engineering intermediate language at roughly the level of GRAIL, OIL or Staab's representation and for a still higher level user-oriented intermediate representation.

A key aspect of *Open*GALEN's intermediate representations is that they are 'soft' and can be adapted to the requirements of individual sites. An intermediate representation consists of a) a set of user oriented 'descriptors' or terms b) mappings of those terms to concepts in the underlying ontology c) a set of constrained templates providing the links between the descriptors d) a set of transformations between the intermediate representation and expressions in the underlying ontology. Within broad limits, sites can author their ontologies using descriptors and templates tailored to their needs and tastes. All intermediate representations can then be transformed into a common underlying representation [21].

However, the transformation process is not infinitely flexible; some consistency is required from the domain authors. Therefore, in addition to the intermediate representation, guidelines and examples are required for semantic normalisation as described below.

Furthermore, although the intermediate representations are relatively comprehensible, in many applications simple generated pseudo-natural language noun phrases are more compact and familiar and can be adapted to the user's own language. *Open*GALEN has found natural language generation essential to user acceptance. Most language generation is general, but additions to lexicons and grammars are usually required for each new application.

Finally, any real application requires a set of quality assurance criteria and the tools to test them. These should be developed and agreed at the same time as the intermediate representations although, in practice, they often evolve in the course of development.

Development phases

*Open*GALEN therefore divides development of large ontologies into two phases: design and population. In practice, these phases are iterative, but it is easier to describe the processes as if they were sequential.

The Design Phase

In the design phase knowledge engineers extends the basic ontological schema and prepares a user-authors' view or Intermediate Representation. The outputs from the design phase configure two sets of tools, one for the knowledge engineers and one for the domain experts. In the population phase, domain experts populate ontology with domain concepts using intermediate representations which is transformed into the common underlying description logic representation.

The goal of the design phase is to produce five related outputs and incorporate them into a set of tools for the domain experts to use in the population phase as shown in Figure 1:

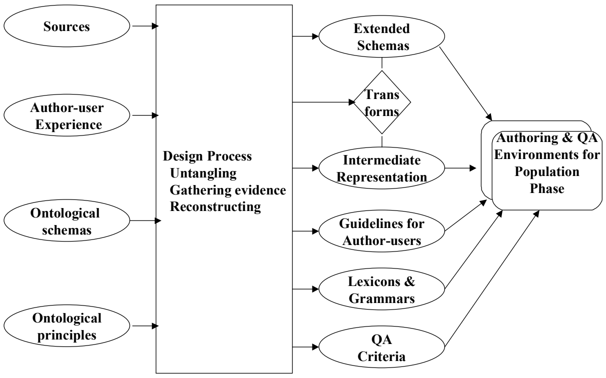

Figure 1: The design phase of OpenGALEN development

- *Intermediate Representations* adapted to each major group of domain expert authors' requirements
- *Guidelines* for domain expert authors
- *Schemas* for the underlying ontology, along with transformation rules from the Intermediate Representation to the underlying ontology.
- *Lexicons and Grammars* for natural language generation for display of results to users
- *Quality assuranc*e(QA) criteria based on the combination of the above three. Ideally quality assurance criteria are set at the time of the original design and modified iteratively, although this ideal is not always achieved in practice.

The four outputs feed into the tools and environments which are used in the population phase.

The population phase

The population phase is best described by two views: a layered view as in Figure 2 showing the different components and how they interact, and an iterative view as in figure 3 showing the flow of information and interaction between domain experts and the central knowledge engineering team.

In the population phase, domain experts usually work from sources such as existing terminologies or classifications. The first step is to paraphrase the phrases or 'rubrics' from those sources into unambiguous statements to be represented in the intermediate representation. Separating the paraphrase step from the representation step allows quality assurance and discussion of the domain experts' interpretation of the sources to be separated from their representation of those sources in an intermediate representation.

To transform the paraphrase into the intermediate representation and organise the results, the domain experts interact with documentation and tools incorporating the quartet of resources developed in the design phase: guidelines, tools, intermediate representations, and quality assurance criteria. This quartet of resources is linked to the formal ontology through the transforms between the intermediate representation and the underlying formal ontology. (The addition of a classifier at the level of FaCT is likely to produce a further layer of Implementation Logic as shown in grey.)

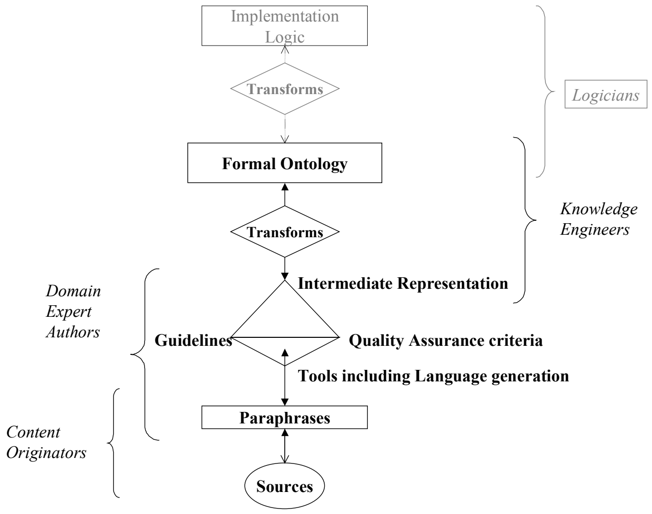

Figure 2: Layering of the population phase of *Open*GALEN development with extrapolation to an additional layer of implementation logic

This methodology provides clear regions of interaction between the various groups involved in the process. The domain expert authors and the knowledge engineers interact over the intermediate representation; the knowledge engineers and logicians over the transformation to the implementation logic and the formal ontology language; the domain experts and the content originators over the paraphrase of the original sources.

In practice the primary interaction is between the domain experts, who often work in independent units, and the knowledge engineers, who are usually a central resource. In general, the domain experts can author content in the intermediate representation and have it structured, classified, and ready for quality assurance without any intervention from the central knowledge engineers. New concepts in existing categories can be authored locally and reported to the centre automatically. However, novel concepts and constructs require the intervention of the central knowledge engineering team as does the overall reconciliation and integration of the work of the centres. This combination of local autonomy and with central support and integration produces the double cycle is shown in Figure 3 which is characteristic of development using GALEN. This pattern is effective at maximising local autonomy and minimising the requirement for central coordination. Overall, in an established area of development, central services require about 10% of the total effort.

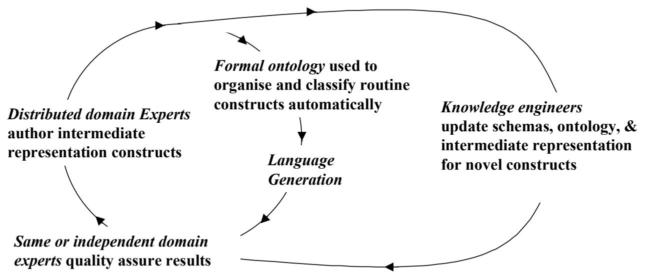

Figure 3: *Open*GALEN Population Cycle

DESIGN ISSUES

The layered architecture has allowed *Open*GALEN to evolved a principled and systematic approach to designing ontologies for clinical applications which addresses a wide range of issues:

- Issues in design of the ontology itself and the transforms between the intermediate representation and ontology. The goal is that domain experts be largely unaware of these choices because they are handled by the transformation between the intermediate representations and the ontology. In most cases, technical changes to the ontology should not require any recourse to the domain experts except, possibly, for additional quality assurance.
- Issues in dealing with the idiosyncrasies of the domain which can only be implemented as guidelines to the domain experts or constraints within the tools and environment.

Key issues from both sets are discussed below along with how their interrelation with the underlying ontology and the intermediate representation.

Issues in the design of the ontology and transforms

Untangling taxonomies

GALEN's source material typically consists of seriously tangled hierarchies, typically derived from 'broader than'/ 'narrower than' constructs in traditional library science and thesauri rather than the formal inferential meaning of subsumption in description logics. The hierarchies typically mix the notions of kinds, parts, function, use etc. The patterns are familiar to users, make for easy access to terms, but make formal inference all but impossible. For example, heart diseases are found in thirteen of the eighteen chapters of the International Classification of Diseases.

GALEN's approach is to separate out each 'axis' into a separate taxonomy of elementary concepts, and then recombine as expressions in the description logic. Where two axes are highly correlated, this can involve introducing much seemingly redundant information – *e.g.* separating the 'action' and 'use' of drugs may lead to recording separately an action of 'bronchodilation' and a indicated for 'bronchodilation'. However, in other cases the use and action may be quite different – *e.g.* an action of 'vasodilation' and a use of 'management of hypertension'.

Operationally, *Open*GALEN maintains the principle modularity by specifying that *elementary* concepts should break down into disjoint taxonomies, *i.e.* each elementary concept should have only one elementary parent and be disjoint from all its 'sibling' elementary concepts. The taxonomies of elementary first class concepts are open – *i.e.* at each level of the hierarchy siblings are disjoint but do not exhaust the parent concept. This reflects the reality that lists of diseases, abnormalities, and even anatomy can almost never be fully exhaustive, especially when the possibility of congenital abnormalities are taken into account. By contrast taxonomies of modifying concepts such as 'severity' may exhaustive and therefore closed. All multiple classification and overlapping of concepts are the result of definitions and descriptions. This may involve creating artefactual concepts known as 'roles', *e.g.* "doctor" is defined as a "person who plays a 'doctor role'" and a hormone as a "substance which plays a 'hormonal role'". (This use of word "role" is not to be confused with the word "Role" for semantic relation in description logic parlance.) This allows clean disjoint taxonomies for the notions of 'organism', 'person', etc. and for 'social role', 'clinical role', 'doctor role', 'patient role', etc.

The structures which result of untangling taxonomies and recombining them through logical definitions are is consistent but contain much detail which is irrelevant to users. Some of this detail can be hidden by using the definitions in the formal ontology itself, but an important function of the intermediate representation is to provide macros to hide the rest.

Locations, parts, wholes & related spatial notions

Much of the power of *Open*GALEN's ontology stems from its distinction between different sorts of part-whole and other spatial relations. Although adapted roughly from Winston's structure[13, 29], it differs from it significantly and distinguishes:

Location – Lesions and abnormalities are 'located' in things rather than part of them or contained in them. (If physical containment is implied, as in foreign bodies it is specified additionally and separately

Parts – in four main flavours

Division – Roughly self similar parts having the same layers, e.g. hand and arm

Layers – horizontal layers such as the skin which extend across divisions

Structural Components – discrete parts which normally reside in only one division

Functional Components – parts of a functional unit which may or may not be contained in or contiguous with the whole, e.g. the various glands which make up the endocrine system.

Containment – physical containment of one structure by another where there is quite different function and origin, e.g. bone marrow in bones

Connections – which may or may not be considered part of the things connected.

There are also distinctions drawn between two-dimensional and three-dimensional parts analogous to those in the Digital Anatomist project [18].

Establishing the original anatomical structures requires careful consideration by the knowledge engineers. Once the anatomical structure is established, the use of the different relations in descriptions of surgical procedures, diseases, etc. can be determined automatically by the transforms between the intermediate representations and the underlying ontology.

The part-whole structure requires the use of axioms similar to Cyc's TRANSFERS-THRO [10] to cope with the paradigm that "diseases/procedures of a parts are diseases/procedures of wholes". The general schema required is equivalent to:

$$R_1 \circ R_2 \rightarrow R_1$$

for at least a restricted set of roles R_1 and R_2. Classification algorithms for description logics supporting such schemas remain an outstanding problem in description logics [2, 5]. GRAIL implements a partial solution for restricted cases [15]. Although recent results concerning fully general solutions are not encouraging [27], practical solutions at least for restricted cases are essential for any effective medical ontology.

Defaults and 'Extrinsics'

A major function of Frame systems is to deal with default knowledge – *i.e.* information which is true in general but subject to exceptions. Formal description logics do not support default reasoning. However, they can provide a framework for separate default reasoners.

In a static system it is always possible to 'compile out' default knowledge out by re-representing each item at all highest levels below which there are no exceptions. However, this strategy is inappropriate for knowledge acquisition and for use in many dynamic systems because it may not provide default values for new information when added which is a key function in many such applications. For example, drug-drug interactions are best specified at the level of drug classes with the exceptions enumerated explicitly, so that when a new drug is added, it acquires the default 'safe' set of interactions unless they are explicitly overridden.

GRAIL provides a special mechanism for attaching 'extrinsic' information to concepts which do not affect their classification but which can be manipulated by a special set of operations based around the notion of retrieving the set of 'most specific' extrinsic information of a given type. This mechanism is also used for handling complex mappings to external classifications and terminologies and for links to natural language applications.

The distinction between extrinsic (default) and intrinsic (definitional and descriptive) information is not at all intuitive to domain experts. The decision as to which constructs should be 'extrinsic' is made by the knowledge engineers and implemented in the transforms between the intermediate representation and description logic so as to be transparent to domain experts.

Reification of relations and 'wrapping'

As stated in Section 2.1, the criteria for correctness in *Open*GALEN is consistent classification and re-use rather than any notion of 'naturalness'. To achieve consistent re-usable representation within the underlying ontology requires a number of complex constructs which are concealed from users by the intermediate representation

- For many purposes, all diseases are 'wrapped' and represented as collections of one or more disease concepts in order to cope with common constructs as used in medical records and existing coding systems such as 'A with B', 'A without B', etc.
- To cope with the fact that GRAIL does not handle negation explicitly and to make the distinction between absence of information and negative information unambiguous, disease and procedures are usually expressed with a second layer of wrapping as 'presence/absence of A', 'Performance/Nonperformance of procedure', etc.
- All modifying relations are reified as 'features' which may be chained in order to allow consistent re-usable patterns. For example "elevated temperature" is represented in the ontology itself analogously to *patient-hasFeature-temperature-hasFeauture-elevation-hasState-elevated* rather than *patient-hasTemperature-elevated.*

All of these transformations are hidden from the domain expert, so that a simple notion which appears to the user as 'Diabetes hasState severe' in the intermediate representation is transformed into an internal representation in the ontology analogous to *ClinicalSituation-involves-(Presence-isExistentialStateOf-(Diabetes-hasFeature-(Severity-hasValue-severe))).*

Dealing with the idiosyncrasies of the domain knowledge

Semantic normalisation

It is easy to agree that all surgical procedure are constituted by an 'act' on some 'thing' which either is, or is located in, an anatomical structure. It is less easy to agree on what constitutes an 'act' when there is a hierarchy of motivations: for example, "inserting a pins to fixate a fractured bone" or "destruction of a polyp by cautery" and "removal of a polyp (by excision)". Furthermore, important classifications hang notions of motivation such as "palliative surgery" and "corrective surgery". In addition, some systems wish to be able to record operations just as 'correction of X' without describing the exact 'act' while others wish to record 'insertion of pins in fractured bone' without recording that the purpose is fixation.

To address this problem, one of the members of the project proposed a classification into four levels: L4 Clinical Goal (palliation, Cure); L3 Physiologic goal: (correction, destruction, …); L2 primary surgical method (excision, insertion, lysis,…); and L1: low level surgical act (cutting, cautery, …) [19]. It is tempting to believe that a consistent list of concepts in each category can be agreed, so that resolution could be done completely automatically. However, intuitions and requirements clashed sufficiently to make this difficult. For example, ‘Cautery’ can sometimes be a mere low level act or sometimes the primary method. These ambiguities are dealt with in the formal ontology by having concepts for “simple cautery” and “removal by cauterisation”.

Concealing such distinctions from the domain experts completely sometimes adds more confusion than it avoids. Therefore, semantic normalisation must be dealt with by a combination of guidelines for how things should be done, and transforms which recognise anomalies, and quality assurance procedures which catch any remaining inconsistencies.

Dealing with implied and normative knowledge

A key problem in dealing with pre-existing terminologies is that much of the information is implied rather stated. Hence the requirement that each term or ‘rubric’ be paraphrased before being represented. Many of the guidelines concern paraphrasing of different sorts of rubrics.

A key part of this process is expanding expressions such as “insertion of pins in the Femur” to “Fixation of Femur by means of insertion of pins”. That the intended meaning includes fixation can only be inferred from context and general medical knowledge – if the insertion were for any other purpose it would be stated in the rubric since fixation constitutes the overwhelming majority of reasons for inserting pins into femurs - so much so that it is not stated in the rubric. It is part of the meaning in context but not of the literal meaning. Similarly many disease classifications depend on normative anatomy which is not invariably true. For example, the thyroid gland is almost always located in the neck but may be ectopically located in the chest.

Idiomatic meaning vs logical definition

A closely related problem occurs when describing important abstractions such as ‘Heart Valve’ or ‘Endocrine Surgery’. These concepts might naturally be defined as “Valve in the heart” or “Surgery on an endocrine organ” respectively. Both produce results which surprise clinicians. ‘Heart Valve’ conventionally means one of the four main valves at the entrance and exit of the ventricles rather than any of the other valvular structures, many of which are normally active only prior to birth. Similarly, ‘Endocrine surgery’ typically refers to a particular set of operations on endocrine organs excluding the reproductive tract, even though all would agree that the gonads have endocrine function.

Achieving a familiar structure: Tagging vs Mapping to original sources

The untangling process by itself provides typically provides descriptions of leaf concepts, but may not provide the higher level abstractions users expect. One important requirement is often to provide additional tagging related to the familiar hierarchies, so that users can find concepts where they expect them. This is usually done by adding the mapping as part of the formal description of low level concepts, and then creating concepts such as ‘drugs in chapter three’ as abstractions.

Note that this tagging is used only to mark high level constructs in the source classifications. Detailed mapping of leaf nodes in source classifications almost always require taking into account special rules of usage unique to that classification and so requires further inference mechanisms beyond the scope of this paper.

Pathology and abnormality

Being ‘normal’ or ‘abnormal’, ‘pathological’ or ‘physiological’ are key notions in medicine. However, what is meant by such terms is a thorny issue in general medical usage let alone formal ontologies. *Open*GALEN has established a consistent approach: “Abnormal” indicates “clinically noteworthy”; “pathological indicates “in need of clinical management (possibly by doing nothing)”. This approach is reflected at all three levels, the ontology, intermediate representations, and guidelines, but the complex interrelations and inferences are confined to the underlying ontology and hidden from domain experts.

AN OUTLINE EXAMPLE

The following is an abbreviated example of the process as applied to analysis of surgical procedure terminjologies from original rubric through paraphrase, intermediate representation, transformation to generated natural language.

RUBRIC: *Insertion of pins in neck of femur*

PARAPHRASE: *Fixation of femur by insertion of pins in neck of femur*

INTERMEDIATE REPRESENTATION:

```
MAIN fixation
  ACTS-ON femur
  BY-MEANS-OF insertion
    ACTS-ON pins
      INTO neck
        IS-PART-OF FEMUR
```

GRAIL/GALEN Ontology

```
Performance which isOf
  ('SurgicalFixation' which
    <actsOn Femur
      hasSubprocedure (Performance which isOf
        'SurgicalInsertion' which
          <actsOn Pins
            hasLocation (AnatomjicalNeck which
                          isLinearDivisionOf Femur>))>)
```

GENERATED LANGUAGE:

“Fixation of femur by means of insertion of pins in neck of femur”

(In the GRAIL representation, concepts in single quotes are further defined elsewhere and which is a keyword introducing a series of attribute value pairs bracketed by <…>. See [15] for full details of notation.)

Note that the transform from intermediate representation to the GALEN Ontology has supplied the context specific mapping of *INTO* to *hasLocation* and *IS-PART-OF* to *IsLinearDivisionOf* based on the classification of *Pins, AnatomicalNeck,* and *Femur.* Given different categories of object and value, *INTO* might have been transformed as *contains* and *IS-PART-OF* might have been transformed as *isComponentOf*, *isLayerOf*, *contains*, etc. Note also that the 'wrapping' *Performance* has been provided which allows for combinations which involve *NonPerformance.* Here the generated language is close to the original paraphrase, but more complex cases lead to less felicitous language.

RESULTS AND DISCUSSIONS

*Open*GALEN has been used in two major areas:

- Developing and maintaining surgical procedure classification in several European countries including being the primary development vehicle for the new classification of surgical procedures in France reconciling previously separate systems used in public and private sectors.
- Developing a drug ontology for use in prescribing support in the United Kingdom as part of the PRODIGY project [9, 14].

The methodology for distributed loosely coupled development has been used primarily in the surgical procedure development during the EU funded GALEN-IN-USE project where it allowed nine centres in seven countries to co-operate on various aspects of developing and integrating surgical procedure classifications. The introduction of the intermediate representation has reduced the training time required for domain experts to roughly three days plus telephone and email support, sometimes supplemented by a one or two further days of advanced training. This contrasts the several months training required for a knowledge engineer to be able to use the underlying ontology. Just as important, it dramatically reduced the time and effort required to reach consensus. Domain experts did not have to deal with what they regarded as arcane distinctions, and controversial decisions could be deferred until sufficient data was gathered to make them based on evidence rather than dogma. The fear of wasting effort because was much reduced, because it was almost always possible to preserve the intermediate representations and change only the transforms to the underlying description logic.

Natural language generation has proved unexpectedly to be essential both to acceptance by users – regardless of how intuitive the intermediate representation appears to designers, simple noun phrases are both more compact and more accessible to domain experts, especially for quality assurance.

At this level, the method appears cost effective compared with the alternative manual development of classifications. Replication is required, but a preliminary study by the Dutch collaborators indicated that the cost of using *Open*GALEN techniques was on the order of 25% that of using conventional techniques even including the one-time-only cost of take on, primarily because the techniques reduced the number of costly meetings of expert committees and led to more rapid consensus [32].

The importance of the approach to 'untangling taxonomies' can perhaps best be illustrated by recent experience with the major medical standards body Health Level Seven (HL7). The seemingly simple problem of classifying the forms and routes by which a medication can be given – "oral tablets", "nasal sprays", "ointments to be rubbed on the skin" had caused serious difficulties. There are at least five different axes involved. Between the various providers of drug information there are over 800 concepts – a fraction of what are involved in other *Open*GALEN knowledge bases but nonetheless, a significant number. Developing a classification manually had proved a daunting task and was still incomplete after over a year; developing the classification using *Open*GALEN's formal methods was completed with a few weeks effort with contributions from five sources, none of them with previous involvement with *Open*GALEN or any formal training [30].

The separation into of the development into design and population phase, and the separation of the design issues between those involving the underlying ontology and those involving the domain itself have improved the ability to reach consensus and vastly reduced the number of arguments – a major cost in ontology development in our experience. This approach contrasts sharply with the more centralised approach taken in the Convergent Terms Project and SNOMED-Reference-Terminology projects [4, 22].

The layered architecture seems to use almost inevitable for the design of large re-usable ontologies. The predecessor application, PEN&PAD [6, 12] based the implementation directly on the ontology without an intervening layer. As a result, the developers frequently succumbed to the temptation to change the ontology to fit the application, sacrificing re-use to expediency. GALEN's intermediate representation provide a simpler alternative.

We believe the need for such intermediate representations will become more, rather than less, critical as more powerful description logics such as FaCT and ShiQ [8] come into use. While it is tempting to believe that OIL [1] will provide a suitable vehicle for direct development, our preliminary experience suggests that it is best treated as a language of similar level to GRAIL – a better vehicle for knowledge engineers but best hidden from domain experts who will still require environments oriented to their specific needs and packaged together with a variety of tools for reasoning, access to information, calculation, and specialist services, presented at a level which corresponds to the issues which concern them, with contact only where necessary with the logical implementation.

At the same time the use of intermediate representations presents an important route for adapting re-usable ontologies to specific applications. For architectures such as PROTÉGÉ, in which specific applications are developed on the basis of an ontology, the hope is that the 'meta authoring' of a suitable intermediate representation and view onto a more general re-usable ontology might replace the repeated development of bespoke ontologies. This is already occurring to some degree within the PRODIGY project [14,20].

ACKNOWLEDGEMENTS

This work has been supported in part by the European Commission and the UK Department of Health.

REFERENCES

1. The OIL Home Page, (2000), www.ontoknowledge.org/oil/
2. Artale, A., E. Franconi, and N. Guarino. Open problems for part-whole relations. in International Workshop on Description Logics (1996) http://www.cs.man.ac.uk/franconi-bin/retrieve/links/franconi-papers/part-papers/part-dl96.dvi.gz
3. Boose, J. and B. Gaines, eds. Knowledge Acquisition Tools for Expert Systems. , ed. Seri es Boose, J. and B. Gaines, Academic Press: London. (1988).
4. Campbell, K., Scalable methodologies for distributed development of logic-based convergent medical terminology. Methods of Information in Medicine, . 37(426-439): .(1998)
5. Franconi, E., Personal Communication, . 2001.
6. Horan, B., *et al.* Supporting a Humanly Impossible Task: The Clinical Human-Computer Environment,. in Interact 90. Elsevier Science Publishers, B.V.North-Holland (1990) 247-252
7. Horrocks, I. Using an expressive description logic: FaCT or Fiction. in Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference on Knowledge Representation (KR 98). Morgan Kaufmann (1998) 634-647
8. Horrocks, I., U. Staller, and S. Tobies, Practical reasoning for very expressive description logics. Journal of the Interest Group in Pure and Applied Logics (IGPL), . 8(3): .(2000) 293-323.
9. Johnson, P.D., *et al.*, Using scenarios in chronic disease management guidelines for primary care. JAMIA, (Symposium Special Issue): .(2000) 389-393.
10. Lenat, D.B. and R.V. Guha, Building Large Knowledge-Based Systems: Representation and inference in the Cyc Project. , Reading, MA: Addison-Wesley. (1989) 372.
11. Musen, M., Modern architectures for intelligent systems: reusable ontologies and problem-solving methods. Journal of the American Medical Informatics Association, (Symposium supplement): .(1998) 46-54.
12. Nowlan, W.A., Clinical workstation: Identifying clinical requirements and understanding clinical information. International Journal of Bio-Medical Computing, . 34: .(1994) 85-94.
13. Odell, J.J., Six different kinds of composition. Journal of Object Oriented Programming, . 5(8): .(1994) 10-15.
14. Purves, I., *et al.* An aristocrat, a precocious child and a sacred cow. Developing general practice clinical computer systems: a collaboration between general practitioners, system suppliers and the NHS Executive - lessons from the Prodigy methodology. in Annual Conference of the Primary Health Care Subgroup of the British Computer Society (1996) 51-73
15. Rector, A., *et al.*, The GRAIL concept modelling language for medical terminology. Artificial Intelligence in Medicine, . 9: .(1997) 139-171.
16. Rector, A., *et al.* Making sound re-usable terminology practical: The GALEN approach. in Towards and Electronic Health Care Record, Europe. Medical Records Institute, Newton Mass (1998) 298-305
17. Rector, A.L., *et al.*, Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine. IEEE Transactions on Information Technology in BioMedicine, . 2(4): .(1999) 229-242.
18. Rosse, C., I.G. Shapiro, and J.F. Brinkley, The Digital Anatomist foundational model: Principles for defining and structuring its concept domain. Journal of the American Medical Informatics Association, (1998 Fall Symposium Special issue): .(1998)
19. Rossi Mori, A., *et al.* An ontological analysis of surgical deeds. in Artificial Intelligence in Medicine Europe (AIME-97) (1997)
20. Solomon, D.S., *et al.*, A reference terminology for drugs. Journal of the American Medical Informatics Assciation, (1999 Special Conference Issue): .(1999)
21. Solomon, W., *et al.*, Having our cake and eating it too: How the GALEN Intermediate Representation reconciles internal complexity with users' requirements for appropriateness and simplicity. Journal of the American Medical Informatics Association, (Fall Symposium Special Issue): .(2000) 819-823.
22. Spackman, K.A., K.E. Campbell, and R.A. Côté, SNOMED-RT: A reference Terminology for Health Care. Journal of the American Medical Informatics Association (JAMIA), ((Symposium special issue)): .(1997) 640-644.
23. Staab, S. and A. Maedche. Ontology engineering beyond the modeling of concepts and relations. in ECAI 2000. 14th European Conference on Artificial Intelligence; Workshop on Applications of Ontologies and Problem-Solving Methods (2000)
24. Tu, S., *et al.*, Ontology-based configuration of problem-solving methods and generation of knowledge-acquisition tools - application of PROTEGE-II to protocol-based decision-support. Artificial Intelligence in Medicine, . 7: .(1995) 257-289.
25. Tu, S.W. and M.A. Musen. A flexible approach to guideline modelling. in AMIA Fall Symposium. Hanley and Belfus (1999) 420-424
26. Uschold, M. and M. Gruninger, Ontologies: principles, methods and applications. Knowledge Engineering Review, . 11(2): .(1996)
27. Wessel, M., Obstacles on the way to spatial reasoning with description logics: undecidability of ALC_{RA}, . 2000, Fachbereich Informatik der Universitat Hamburg. p. 27.
28. Wielinga, B., *et al.* The KADS Knowledge Modelling Approach. in The Japanese Knowledge Acquisition Workshop (JKAW '92) (1992)
29. Winston, M., R. Chaffin, and D. Hermann, A taxonomy of part-whole relations. Cognitive Science, . 11: .(1987) 417-444.
30. Wroe, C. and J. Cimino. Using openGALEN techniques to develop the HL7 drug formulation vocabulary. in American Medical Informatics Association Fall Symposium(AMIA-2001) (2001) (Submitted for publication)
31. Wroe, C., *et al.*, Inheritance of drug information. Journal of the American Medical Informatics Association, (Annual Symposium Special Issue): .(2000) (in press).
32. Zanstra, P., Personal communication, . 1999.