Having our cake and eating it too: How the GALEN Intermediate Representation reconciles internal complexity with users' requirements for appropriateness and simplicity

W. D. Solomon[1], A. Roberts[1], J. E. Rogers[1], C. J. Wroe[1], A. L. Rector[1]
[1]Medical Informatics Group,
Department of Computer Science, University of Manchester, UK
& OpenGALEN www.opengalen.org

Clinical terminologies are complex objects, getting more complex as the requirements on them grow, and as more complex technologies are used in their construction. But to the clinical end-user, functionality and utility is important, not inherent complexity — the simpler a clinical terminology can be for the end-user, the better.

To reconcile these contradictory requirements, the GALEN Programme has developed an Intermediate Representation that allows the OpenGALEN Clinical Terminology to retain a high degree of internal complexity, whilst allowing it to be efficiently maintained, and easily used. This paper describes the elements of the Intermediate Representation, how it works, and some experience of its use.

INTRODUCTION

Clinical terminologies have contradictory requirements. On the one hand, they need to be complex to represent medical concepts. On the other, users require simplicity. The complexity of clinical terminologies has several axis:

1) The scale of medicine forces them to be large.

2) As the requirements on clinical terminologies grow, so their complexity grows.

3) New levels of detail and technologies are being brought to bear on the issue. For example, NHS CTV3 includes semantic descriptions of the structure of terms within it[1]. SNOMED RT takes this a step further and uses an external description logic classifier to help organise its structure[2].

The GALEN Programme[3, 4] also uses description logics, and is, by necessity, complex; the *Open*GALEN Clinical Terminology[5] is designed to support many applications. It can be delivered as an active resource within a GALEN Terminology Server[6], as well as forming part of the support infrastructure for the maintenance and delivery of other coding and classification systems[7].

Description logics are complex

The experience of the GALEN Programme has been that the use of description logics brings its own problems. The level of detail and structure that description logics allow - and which are needed if the clinical terminology is to live up to its new expectations - makes authoring and maintenance more difficult. See for example Rogers and Rector[8], which illustrates the range of partonomic relationships now required. Building and maintaining a clinical terminology in such an environment requires more than traditional clinical skills.

Users require simplicity

A successful clinical terminology must disappear into the underlying infrastructure; the end-user simply wants better tools with which to do a better job.

This applies equally whether the end-user is a physician entering documentation into an electronic patient record, or producing a clinical terminology.

The GALEN Intermediate Representation

To reconcile these contradictory requirements, we have developed a layer of abstraction between clinical end-users and the underlying technology: the *GALEN Intermediate Representation (GALEN IR).* This acts as a high-level language for GRAIL, the description logic used in the *Open*GALEN Clinical Terminology[9]. We have previously described motivation for the GALEN IR in detail[10]. This paper describes for the first time the elements of the GALEN IR, how it works, and experiences of its use.

REQUIREMENTS ON THE GALEN IR

The GALEN IR has been designed to satisfy the following requirements.

1) Allow clinicians to use clinical knowledge to build the terminology without specialised knowledge of description logics or the style of the underlying model. This overcomes a knowledge acquisition bottleneck and optimises use of clinical knowledge.

2) Separation of clinical knowledge authoring from detailed internal implementation choices.

3) We must be able to build efficient software to translate between the GALEN IR and GRAIL.

4) Definition of simpler interfaces to the *Open*GALEN Clinical Terminology for other systems (see for example Webber et al[11]).

5) We must be able to build development environments to support its use.

6) Support for providing guidance for end-users.

There are also a number of requirements specific to the GRAIL target language:

- The use of words closer to day-to-day clinical language than GRAIL knowledge-names.
- The ability to use different natural languages
- The use of a small number of semantic links, instead of the large number of GRAIL attributes.
- The use of a simple recursive structure rather than the complex, multi-attribute-concept chains that are necessary in GRAIL.
- The ability to say the same thing in different ways, allowing for local variation in style.

SATISFYING THE REQUIREMENTS

An example

We illustrate the GALEN IR using the surgical procedure that excises a lobe of the lung. The 'dissection', or semantic description of this concept in GALEN IR, is shown in Figure 1.

```
RUBRIC "excision of lobe of lung"
MAIN excising
    ACTS_ON lobe
        IS_PART_OF lung
```

Figure 1: an example dissection

Words in lowercase, designed to be familiar to any clinician, are known as 'descriptors'. Words in uppercase are 'links', and users of the GALEN IR must be able to appropriately use a small number of them. Indentation represents recursive structure: the lobe of the lung is what is excised. Figure 2 shows the underlying GRAIL representation for the concept.

```
(SurgicalDeed which
    isMainlyCharacterisedBy (performance whichG
        isEnactmentOf (Excising which
            playsClinicalRole SurgicalRole) which
                actsSpecificallyOn (Lobe which
                    isSpecificSolidDivisionOf Lung))).
```

Figure 2: GRAIL expression representing 'excision of lobe of lung'

Overview of the GALEN IR

The GALEN IR consists of a number of components:

- A controlled vocabulary of descriptors and links that are mapped to elements of the *Open*GALEN Clinical Terminology, and are sufficient so as to be able to describe clinical concepts.
- Tools and environments to build and maintain that resource.
- A specification of how structures containing elements are automatically translated into GRAIL concept expressions;
- Implementations of that specification. (Two implementations have been illustrated in the GALEN programme).

The overall flow, with all these components, is shown in Figure 3:

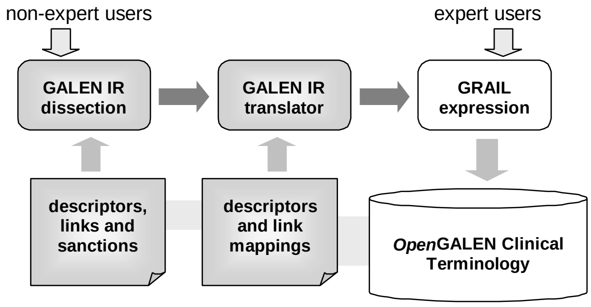

Figure 3: GALEN IR components and interactions

Without the GALEN IR the only entry route to the *Open*GALEN Clinical Terminology is for expert users, using GRAIL. With the GALEN IR, an extra entry-point for non-expert users is provided. It becomes possible to build end-user tools (such as the Classification Workbench[7]) requiring less specialised knowledge of the internal description logic.

DISSECTIONS TO GRAIL

The translation of dissections to GRAIL makes use of three kinds of knowledge: descriptor mappings, link mappings and structural transformations. It follows the following steps:

- Rearrangement and transformation
- Translation of the dissection start point
- Recursive translation of the dissection body
    - Descriptor to concept mapping
    - Descriptor mapping rules
    - Link to attribute mapping

The translation is described below. Examples of knowledge used in the translation are shown as abstract rules in Figure 4, and referred to throughout the description. The complete translation of the dissection in Figure 1 to the GRAIL in Figure 2 is shown in Figure 5.

Rearrangement and transformation

For modelling brevity, a single dissection can represent several similar concepts. These are folded in to a single dissection using various disjunction keywords. Before translation of individual dissections, these are unfolded in to sets of dissections. An example is shown in Figure 4a. Other transformation rules can be applied – for example all of a dissection after a given link can be suppressed (useful for partial compilations), shown in Figure 4b.

Start point

Each GRAIL concept starts with a wrapper — in Figure 2, SurgicalDeed. The GALEN IR allows the definition of separate *configurations*, to which each dissection belongs; for example, the configuration for surgical procedures knows that the start point wrapper is the GRAIL concept SurgicalDeed. This is then treated as if it were an invisible object placed before the link 'MAIN', as shown in Figure 4c.

Recursively translate the dissection body

The remainder of the translation is a simple recursive walk: for each descriptor-link-descriptor triple (starting with the invisible start-point concept described earlier) translate descriptors into GRAIL concepts, apply descriptor mapping rules, and then translate the link. These steps are described below.

(a) Dissection expansion rule

MAIN x_1 OR_MAIN x_2 ->
{MAIN x_1, MAIN x_2}

(b) Link suppression rules

If l:S then d_1 l d_2 -> d_1

(c) Adding a start point

MAIN x -> c_{start} MAIN x

(d) Descriptor to concept mapping

l? d -> l? c

(e) Specialisation rule

if c:C_s and l:L
then l c -> l c which $a_s c_s$

(f) Wrapping rule

if c:C_s and l:L
then l c -> l c_w which a_w c

(g) Chains of attributes

l -> which a c_1…
…which $a_{n-1}c_{n-1}$ which a_n x

(h) Context-specific link disambiguation

c_t:C_t l c_v:C_v -> c_t {A_1 … A_n} c_v

(i) Sanctioning information in link disambiguation

For each possible link to attribute chain mapping in {A_1 … A_n}, the first attribute chain is chosen where a sanction exists in the OpenGALEN terminology for both c_t a c_1 and c_{n-1} a_n c_v. In the special case where A = a, c_t a c_v is checked.

Key

x: dissection body; **d**: descriptor; **l**: link; **S**: set of suppressed links; **l:S** link of type S; **c**: GRAIL concept; **a**: GRAIL attribute; **c:C** concept of type C; **L**: set of links defined for a specific rule; $\mathbf{a_s c_s}$ and $\mathbf{c_w}$ $\mathbf{a_w}$: pairs of defined concepts and attributes used in specific subsumption and wrapping rules; **A**: attribute chain of the form a c_1…which $a_{n-1}c_{n-1}$ a_n

Figure 4: Example abstract rules for GALEN IR to GRAIL translation

Descriptor to concept mapping

Usually, a descriptor maps to a single GRAIL concept. Occasionally it is useful to be able to provide different descriptor mappings according to the context — in terms of the link — with which it is used. For example, the descriptor 'bone' may have different mappings when describing a piece of anatomy, or graft material. The general mapping rule is shown in Figure 4d.

Many descriptor mappings are trivial: for example in the mapping of the descriptor 'lobe' to the GRAIL concept Lobe. This is true only given that the language of the descriptor (English) is the same as that used in the GRAIL authoring.

Some examples are not so trivial, either to cope with inherent ambiguity in natural language, or to cope with specifics of the GRAIL description logic. For example, the term 'appendix' is ambiguous, yet within specific contexts of use is well-understood; so the ambiguous descriptor 'appendix' is mapped to the more explicit GRAIL concept AppendixVermiformis.

In many cases, descriptors used by authors will be defined within the GALEN IR. A problem is encountered when authors add new descriptors with no mappings. To solve this, authors assign a category to new descriptors. They are then mapped via this category to roughly the right part of the *Open*GALEN Clinical Terminology. Further accurate mapping is then required by GRAIL modellers.

Descriptor mapping rules

Additional complexity is handled with *descriptor mapping rules*. For example, the GRAIL expression representing the 'deed' part of a surgical procedure needs to be 'surgical excising' rather than just 'excising'. Rather than map the descriptor 'excising' (and other deeds) to this more complex expression, we add a *specialisation rule* stating: "if the descriptor maps to a kind of GenericProcess, and follows one of a specified set of links, add the criterion 'which playsClinicalRole SurgicalRole' ". So the Descriptor "excising" in Figure 1 is mapped to the concept "Excising" which is transformed by a specialisation rule to "Excising which playsClinicalRole SurgicalRole. This is shown in Figure 4e.

We can also add *wrapping rules*. For example, dissections may use a descriptor that maps to an active drug ingredient as shorthand for a drug, So the descriptor 'Atenolol' is read as 'drug with active ingredient atenolol'. The descriptor needs to be translated to the GRAIL 'Drug which hasIngredient Atenolol'. The wrapping rule states: "if the descriptor maps to a kind of ActiveDrugIngredient, and follows one of a specified set of links, wrap with 'Drug which hasIngredient' ". This is shown in Figure 4f.

Another mapping variant makes the GALEN IR more powerful. The granularity of descriptors can be controlled by mapping to other dissections, and thence to GRAIL, rather to GRAIL directly.

Link to attribute mapping

Links map to either a single GRAIL attribute, or to a 'concept-attribute-concept' (or more complex) chain.

The general case is shown in Figure 4g. In the figure 1, the link 'MAIN' is mapped to the chain

isMainlyCharacterisedBy performance
which G isEnactmentOf

This mechanism is also useful for translating features; for example the link 'HAS_SEVERITY' needs to be translated into the GRAIL chain:

hasSeverity Severity
which hasAbsoluteState

Each GALEN IR link is potentially translated into one of many GRAIL attributes or chains. This is particularly true of partonomic relationships, where the *Open*GALEN Clinical Terminology maintains many different flavours of links, for example to cope with the fact that the relationships between a) the shaft of a femur and the femur, and b) between the femur and the leg, have different meanings[8].

To keep matters simple for the users of the GALEN IR, we have just two partonomic links: 'HAS_PART' and 'IS_PART_OF'; this offers a prime example of the mechanism within the GALEN IR for expanding IR links to GRAIL attributes. It is a two-stage process; the first stage requires additional information: context-specific link disambiguation, and the second stage uses sanctioning information already present in the OpenGALEN Clinical Terminology for further disambiguation. This is shown in Figure 4h and 4i.

Context-specific link disambiguation: each link may have multiple potential mappings to a *candidate set* of GRAIL attributes, controlled by the combination of the type of GRAIL concept of the 'topic' of (descriptor before) the link and the 'value' of (descriptor after) the link. The link mapping rules say that, for the context *lobe* - *IS_PART_OF – lung* (Figure 1) the correct translation is by using one of a candidate set of GRAIL attributes including isSpecificSolidDivisionOf whereas in another dissection containing *skin* - *IS_PART_OF* - *leg*

the correct translation uses one of a candidate set that includes the GRAIL attribute isSurfaceDivisionOf.

The authoring of these rules is kept manageable by exploiting the use of the classification structure in the *Open*GALEN Clinical Terminology — the rules are described at as general a level as possible.

Sanctioning information: another axis of disambiguation from the candidate sets as described in the previous section, is provided by the sanctions present in the *Open*GALEN Clinical Terminology: only the GRAIL attribute within a candidate set that is suitably sanctioned will be used as the final choice in the translation of the original IR link.

In the examples, links are translated into 'defining' properties (definitional features) of the resulting GRAIL. In some circumstances, for example when describing the properties of concepts (such as the indications for a drug) link mappings can be annotated with a non-definitional GRAIL property.

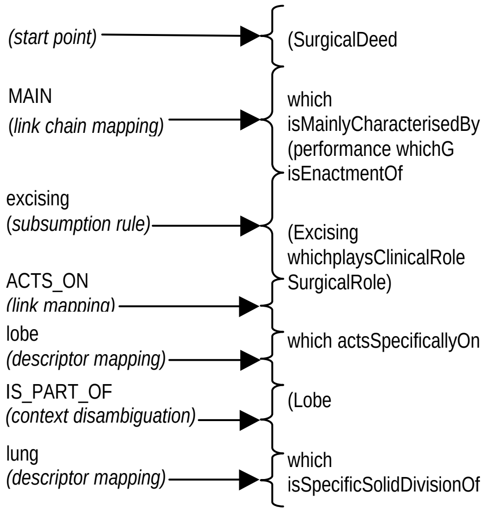

Figure 5: A complete translation (dissection indentation not taken in to account)

EXPERIENCES OF USING THE GALEN IR

The GALEN IR has been used in the GALEN-IN-USE[10] project, and to build a drug ontology for the Prodigy project[12].

Within the Prodigy drug ontology project, the GALEN IR has been used for the rapid modelling of large numbers of concepts by a small number of modellers. It has been most useful for initial knowledge capture of drug information, especially the semi-automated capture of knowledge from semi-structured sources. It has proved useful for adding codes to concepts, providing external references to concepts, and simple compositional modelling. Attempts have also been made to use the GALEN IR as a complete high level language for knowledge authoring, but this has proved unrealistic: the simplicity of the GALEN IR means that it has sacrificed the full expressiveness of GRAIL.

In GALEN-IN-USE, approximately 50 people were trained in nine centres. The usual training period was a three day workshop plus follow-up support. The rate of dissection production varied between centres and the complexity of the rubrics being dissected, but averaged around 50 rubrics per person-day including quality assurance.

Informal evaluation by the Dutch GALEN-IN-USE partners suggests that the effort required to revise their surgical procedure was significantly reduced by use of the GALEN IR, as compared to conventional means. Even though the effort included a significant 'front loading' of initial dissections which would be re-used in future revisions, technical effort was reduced from over five person months to 2.5 person months. Despite this reduction, the number of issues raised and revisions made was increased more than five-fold. Most importantly, the total effort by experts in committees was reduced from 60 person days to 5 person days. These results require replication, but are consistent with the experience of the French partners in developing their new national classification of procedures[13]. This suggests that the GALEN IR can produce a dramatic improvement in productivity in maintaining and revising classifications.

It was anticipated that translation of some dissections to GRAIL might require manual disambiguation. This has so far proved unnecessary: all translation has been fully automatic. The transformation rules plus the constraints provided by GRAIL's sanctioning mechanisms have so far proved adequate.

The descriptor category mechanism, added to cope with unseen descriptors, usually allows reasonable classification for initial checking, but requires accurate GRAIL modelling before classification is fully correct. Within surgery such central modelling effort has averaged on the order of 10% of the effort to produce and quality assure the dissections initially.

CONCLUSION

We have discussed a way to overcome the internal complexity of the *Open*GALEN Clinical Terminology in order to make it easy and efficient to build and maintain. Revisiting the requirements, the use of the GALEN IR has enabled us to:

1. Build tools so that clinicians can contribute with minimal training and additional knowledge[7];
2. Experiment with different ways of translating dissections[14];
3. Build efficient implementations of the dissection translation algorithm;
4. Provide a target representation for other systems[11];
5. Build a development environment to support its use[15];
6. Build tools that contain constraints to guide end-users in their use of the GALEN IR[7].

We believe that the *Open*GALEN Clinical Terminology — and the GRAIL description logic in which it is written — is complex because it needs to be – not because it is over-engineered. To be usable, that complexity has to be managed. We have used a strategy analogous to that found in other fields: create a level of abstraction between the end-user and internal complexity. The complexity is not removed; it is isolated so that it can be managed.

Acknowledgements

The GALEN IR has been developed, refined and used during the GALEN, GALEN-IN-USE and SynEx projects (funded by the EU) and the Prodigy project (funded by the UK Department of Health).

REFERENCES

1. O'Neil M, Payne C, Read J. Read Codes Version 3: A user led terminology. Methods of Information in Medicine 1995;34:187-192.
2. 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.
3. Rector A, Nowlan W, Glowinski A. Goals for Concept Representation in the GALEN project. In: 17th Annual Symposium on Computer Applications in Medical Care (SCAMC-93); 1993: McGraw Hill; 1993. p. 414-418.
4. GALEN. http://www.cs.man.ac.uk/mig/galen.
5. *Open*GALEN. www.opengalen.org.
6. Rector A, Solomon W, Nowlan W, Rush T. A Terminology Server for Medical Language and Medical Information Systems. Methods of Information in Medicine 1995;34:147-157.
7. Rector AL, Baud R, Ceusters W, Claassen W, Rodrigues J-M, Mori JRR, et al. A comprehensive approach to developing and integrating multilingual classifications: GALEN's Classification Workbench. Journal of the American Medical Informatics Association 1998; Fall Symposium Special Issue:1115.
8. Rogers J, Rector A. GALEN's Model of Parts and Wholes: Experience and Comparisons. Accepted for: AMIA Annual Symposium 2000; Los Angeles, California; 2000.
9. 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.
10. Rogers J, Solomon W, 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.
11. Webber B, Markert K, Hardiker N, Rauch B. Towards Consistent Minimal Terminologies. In: IMIA WG 6 Conference; 1999; Phoenix, Arizona, USA; 1999.
12. Solomon DS, Wroe C, Rogers JE, Rector A. A reference terminology for drugs. In: AMIA Fall Symposium 1999; Washington; 1999.
13. Rodrigues, J. M., Trombert-Paviot, B., Baud, R., Wagner, J., Rusch, P. and Meusnier, F. Galen-In-Use: An EU Project applied to the development of a new national coding system for surgical procedures: NCAM. In: Medical Informatics Europe (MIE-97); Thesalonika, Greece: IOS Press; 1997. p. 897-901.
14. Wroe C, Solomon W, Rogers J, Rector A. Inheritance of Drug Information. Poster accepted for: AMIA Annual Symposium 2000; Los Angeles, California; 2000.
15. Solomon W, Rogers J, Rector A, Haring E, Zanstra P. Poster: Supporting the use of the GALEN Intermediate Representation. In: AMIA Fall Symposium 1998; Orlando, Florida; 1998.