

Rubrics to Dissections to GRAIL to Classifications

JE Rogers¹ WD Solomon¹ AL Rector¹ P Pole¹ P Zanstra² E van der Haring²

1) *Medical Informatics Group, Department of Computer Science, University of Manchester, UK*

2) *Medical Informatics Group, Catholic University of Nijmegen, Netherlands*

<http://www.cs.man.ac.uk/mig/galen/> or e-mail galen@cs.man.ac.uk

Abstract. This paper summarises the process in the GALEN-IN-USE project by which rubrics from traditional medical coding schemes are analysed into an intermediate, relatively informal conceptual representation which is then automatically translated into the GRAIL formalism and its Common Reference Model.

1. Introduction

GALEN-IN-USE is an EU funded project, a major goal of which is the development of tools and methods to assist in the collaborative construction and maintenance of surgical procedure classifications. Techniques and tools developed in the previous GALEN project [1,2,3,4,5,6] will support this task. Taking part in the initial phase are four national coding and classification centres: WCC (Netherlands), SPRI (Sweden), CNR (Italy) and University of Ste. Etienne (France). The goal is to author, using the GRAIL formalism [6], conceptual representations of individual surgical procedures, with each centre covering roughly one quarter of the total surgical domain. A combined total of 15-20,000 individual representations will be authored by the end of the project. The resulting GRAIL representations will be integrated into the existing GALEN Common Reference Model (CRM) [1,3,4,5]. This will allow:

- an initial classification of the represented procedures to be automatically derived, based on the knowledge explicitly authored in the analyses and the knowledge already in the reference model;
- machine generation of natural language expressions for all representations, in five European languages;
- refinement, extension and reorganisation of the classification using new classification management tools.

More than 20 individual clinicians have been recruited to analyse original code rubrics into conceptual representations, but most have little or no prior experience of the GRAIL formalism or of the particular ontology and modelling style of the Common Reference Model (CRM). This presented a significant challenge to the project: how to reduce the need for training to occur in the complexities of GRAIL and the CRM before any work could begin.

2. An Intermediate Representation

The solution proposed to the training problem was to begin work using a simpler, intermediate conceptual representation [7]. This was originally conceived as a migration step towards eventually authoring directly in GRAIL. The representation was designed in such a way that conceptual representations authored using it might then be automatically, or semi-automatically, expanded into GRAIL. The representation also allows the authors to capture some concepts which the GRAIL formalism in its present form is unable to handle. Finally, the representation serves as the preferred format in which the centres examine and validate their own, and each other's work. The intermediate

representation is broadly similar to those used by the CANON group or MEDS [8,9,10,11]. It is characterised by:

- a relatively small set of semantic links (ACTS_ON, IS_PART_OF), compared to the CRM;
- a two-tier domain ontology (known as the ‘descriptor list’) specific to the surgical domain. Descriptors (leg, excising, tumour) are explicitly typed by one of a small number of classes (e.g. anatomy, deed, lesion);
- a set of constraints, declaring which links may be used with which descriptor classes;
- a grammar defining a layout, or ‘template’, for well-formed representations.

Domain experts in the centres work using existing local coding schemes (WCC, NCSP etc.) to scope their task. The rubrics (text) and associated codes, but not the original hierarchy, are extracted from the scheme. Working on sections of a few hundred related rubrics at one time, each rubric is manually analysed to author a conceptual representation of its meaning, using the intermediate representation. The immediate result of this analysis and authoring is called a ‘dissection’ of the rubric.

The four centres produced more than 1200 ‘dissections’ in the first four months. Figure 1 shows four completed dissections. Each has a header section, containing information about the original rubric and coding scheme, followed by the conceptual representation itself, introduced by the MAIN keyword. The semantic links are capitalised. Descriptors appear in lower case, preceded by their descriptor class. Initially, authoring involved directly editing an ASCII text file. Any convenience and familiarity which this afforded to the authors was, however, outweighed by the numerous spelling and formatting errors which resulted, preventing satisfactory parsing of the interchange file into the GRAIL expansion environment. Subsequent analyses will be authored using a purpose-built tool, the Surgical Procedure Editing Tool (SPET).

Figure 1.

RUBRIC "operations on papillary muscle" CODE "35.31" MAIN deed:surgical deed ACTS_ON anatomy:papillary muscle HAS_OTHER_FEATURE method VALUE induced arrest of heart	RUBRIC "reattachment of papillary muscle" CODE "35.31.i2" MAIN deed:repairing ACTS_ON anatomy:papillary muscle BY_TECHNIQUE deed:reattaching ACTS_ON anatomy:papillary muscle HAS_OTHER_FEATURE method VALUE induced arrest of heart
RUBRIC "dividing of papillary muscle" CODE "35.31.i1" MAIN deed:dividing ACTS_ON anatomy:papillary muscle HAS_OTHER_FEATURE method VALUE induced arrest of heart	RUBRIC "repair of papillary muscle" CODE "35.31.i3" MAIN deed:repairing ACTS_ON anatomy:papillary muscle HAS_OTHER_FEATURE method VALUE induced arrest of heart

3. Expanding dissections into GRAIL

Dissections authored in the intermediate representation are subsequently imported into an environment (TIGGER) built to manage the process of converting them into GRAIL, simultaneously translating them into both the ontology and the style of the CRM. Imported dissections are first parsed for syntax and for whether they comply with certain agreed modelling conventions (such as that all deeds must ACTS_ON something, a convention taken from CEN [12]). The final GRAIL produced is generally more complex than the dissection from which it comes - sometimes very much more so. The translation process is known as ‘expanding’ and the GRAIL produced from a single dissection as its GRAIL ‘expansion’.

Figure 2 shows a dissection (left) and an automatically generated expansion (right). An expansion comprises a GRAIL representation of the concept at hand, and a series of statements attaching incidental, non-classificatory information to that GRAIL concept, such as the text of the original rubric or the name of the original source file. TIGGER automatically generates, in batches or individually, a GRAIL expansion for each original dissection. However, automatic expansions can only be considered candidate GRAIL conceptual representations of the original rubrics: some may be rejected as invalid when presented to a terminology server. This may occur, for example, if there is a cardinality conflict

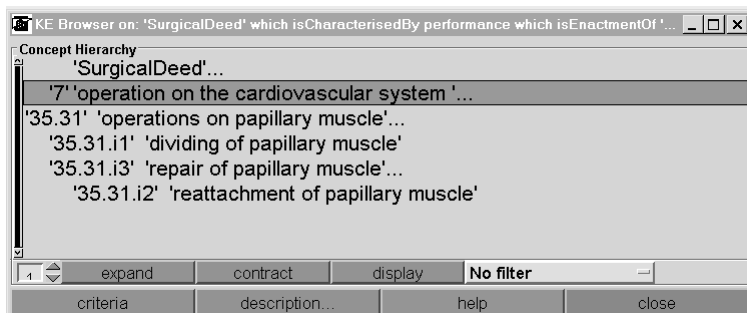
between CRM semantic links used in the expansion. Rejection may indicate a problem which requires alterations to the original dissection, or to the CRM. A few dissections whose expansions are rejected may require their GRAIL representation to be done manually, bypassing the intermediate representation completely.

Figure 2.

Original Dissection	Generated GRAIL expansion
RUBRIC "dividing of papillary muscle" CODE "35.31.i1" MAIN deed:dividing ACTS_ON anatomy:papillary muscle HAS_OTHER_FEATURE method VALUE induced arrest of heart	<pre>(SurgicalDeed which isCharacterisedBy (performance whichG isEnactmentOf (Dividing whichG < playsClinicalRole SurgicalRole actsSpecificallyOn PapillaryMuscle hasSubprocess InducedCardiacArrest>))) extrinsically hasRubric 'dividing of papillary muscle'; extrinsically hasCode '35.31.i1'; extrinsically hasPhysicalSource 'cnr.txt'.</pre>

Sets of automatic expansions produced in this way are presented to the terminology server for classification. The resulting hierarchy of valid expansions may be browsed in a number of ways. The screenshot (figure 3) shows the automatic classification which is derived for the four dissections given in figure 1. The classification of 'operations on the papillary muscle' as a kind of 'operation on the cardiovascular system' occurs because the CRM 'knows' that the papillary muscle is part of the heart which is, in turn, a component of the cardiovascular system. The GRAIL refinement operation [4] is used to declare that actions on part of something are subsumed by actions on the whole. By contrast, the classification of 'reattachment of papillary muscle' as a child of 'repair of papillary muscle' instead of a sibling (as in the original coding scheme) comes from the knowledge explicitly authored in the intermediate representation.

Figure 3.



'repair of papillary muscle' instead of a sibling (as in the original coding scheme) comes from the knowledge explicitly authored in the intermediate representation.

To do the automatic expansion of dissections, however, the TIGGER first requires explicit CRM mappings to be declared for both the descriptors and the links used in a given dissection.

1 Mapping the descriptors

This task of declaring mappings for descriptors is performed by modellers already familiar with the GRAIL formalism and the CRM ontology and style. Once declared, a mapping is presumed valid for all subsequent occurrences of the descriptor in any batch of dissections. As more dissections are processed, the list of already mapped descriptors grows. The accumulated, mapped descriptor list is made available via the SPET to all collaborating authors as a suggested core ontology for use in the next round of authoring dissections.

This approach has the advantage that the ontology with which the dissection authors must become familiar is initially quite small. Further, its growth is owned by the authors themselves but can be guided by those familiar with the GALEN ontology. This contrasts with the already large and complex ontology in the CRM, with which the authors would need to be familiar to author directly in GRAIL. Because the process of getting from dissections to GRAIL is planned to be unidirectional, this methodology also permits some redundancy or duplication of descriptors, as many descriptors can be mapped to the same CRM concept.

Mapping a given descriptor to the CRM is informed by inspection, using TIGGER, of all the dissections which employ it - either in the current batch or in all batches processed so far. Such inspection may provide clarification of what is meant by a descriptor, but may also reveal that one

descriptor has been used with very different intentions by different authors. A mechanism exists for rejecting such ambiguous descriptors and their associated dissections at this stage, to invite re-authoring.

2 Mapping the Links

The links available to dissection authors were chosen such that each is equivalent in intention to the common parent of a range of more expressive links already present in the CRM. Figure 4, for example, shows the CRM links which are to be considered by TIGGER as possible default mappings for the dissection link IS_PART_OF. The mappings from dissection links to the CRM are, therefore, necessarily one-to-many and are declared by the same team of modellers undertaking the descriptor mapping.

To expand a dissection link, TIGGER must determine which, if any, of its candidate CRM link mappings is most appropriate. To achieve this, TIGGER ‘translates’ the descriptors either side of a dissection link into their declared CRM entity mappings. The candidate CRM link mappings are then

Figure 4.

Dissection Link	Default possible CRM mappings
IS_PART_OF	isSolidRegionOf isSpecificStructuralComponentOf isStructuralComponentOf isSpecificSolidDivisionOf isSolidDivisionOf isSpecificLinearDivisionOf isLinearDivisionOf isSpecificSurfaceDivisionOf isSurfaceDivisionOf isSpecificLayerOf isLayerOf

tested in list order: the first one permitted to be used in the CRM between the two entities is chosen. Thus, the dissection fragment:

segment IS_PART_OF intestine
 is expanded into the GRAIL:
 Segment which isSpecificLinearDivisionOf Intestine.
 because the CRM includes the constraint:
 Segment sensibly isSpecificLinearDivisionOf
 TubularBodyStructure.
 This mechanism can also be used to detect dissections which, whilst considered ‘well formed’ within the limited dissection grammar, are still semantically incorrect. For example:
 MAIN excising
 ACTS_ON tumour

IS_PART_OF liver

...can not be expanded: none of the candidate mappings of IS_PART_OF is permitted to be used in the CRM between [Excising], a process, and [Liver], a structure. In future, some semantically incorrect dissections may be rejected at the authoring stage - the SPET will use both the dissection grammar and a limited constraint mechanism to prevent certain link-descriptor pairings being entered at all.

4. Limitations and Future extensions

Coding scheme rubrics frequently contain disjunctions (Excision of tumour or cyst) or conjunctions (Drainage and marsupialisation of cyst). To make the intermediate representation simple, bracketing was omitted from its syntax. As a result, it can only support relatively simple conjunctions or disjunctions. Rubrics with more complex relationships (e.g. partial or complete excision of tibia and fibula, with prosthesis) must be manually enumerated in all their logically and semantically correct conceptual forms. True negation also remains unsupported at present by the GRAIL formalism. However, most of the rubrics which might at first appear to require negation are exclusion or exception criteria. The intermediate representation includes mechanisms for identifying such criteria, which may then be handled within GRAIL by various modelling workarounds or future extensions to the formalism itself.

The relatively relaxed approach to building the descriptor list, with its rudimentary and *ad hoc* class hierarchy, risks it growing to unmanageable and un-navigable proportions. Imposing a more formal organisation would be to some extent to re-invent GRAIL. This may be obviated by offering navigation of the set of mapped descriptors via the hierarchy of their corresponding CRM conceptual mappings.

Both the descriptor list, and the CRM itself, are presently authored in English. This is not the first language of many dissection authors. The generation of natural language expressions for the final GRAIL concepts requires annotations in the destination languages for each CRM primitive. A mechanism to address both of these problems is being studied: dissections will be authored in the local languages, using a local language descriptor list. Each authoring centre will separately maintain a many-to-one translation table from the local terms to a common, English descriptor list shared between all centres. The various translation tables, and the explicit English descriptor to CRM mappings, may be read backwards to derive a list of possible linguistic annotations.

5. Discussion and Conclusion

Although originally conceived as a migration step towards full GRAIL authoring, the intermediate representation and automated expansion process have proved effective in their own right. A high proportion of rubrics can be represented and reliably expanded automatically, and it may be more efficient to author the small remainder directly in GRAIL than to make further enhancements. The success of the intermediate representation is such that, whilst it remains the intention of the project to export CRM modelling expertise from its current localised base, this is no longer on the critical path for the immediate task of building a surgical procedure classification. Further, adding an intermediate layer between knowledge authors and the final representation serves to insulate them from changes in the CRM, and allows those changes to take place more easily and with less disruption.

We are confident that much of the surgical procedure domain can continue to be captured using the intermediate representation, facilitating the involvement of many domain experts by deferring indefinitely any need for them to become familiar with GRAIL or the CRM. A useful by-product of the process is that the act of declaring link and descriptor mappings is building a partial meta-model description of the CRM ontology and style. This will form an important resource for the exporting of CORE modelling expertise, when that occurs.

*With thanks to the centres involved in dissection authoring, and others in the GALEN-IN-USE consortium
GALEN-IN-USE is funded as part Framework IV of the EC Healthcare Telematics research program.*

6. References:

1. Rector, A. (1994). Compositional models of medical concepts: towards re-usable application-independent medical terminologies. in Knowledge and Decisions in Health Telematics P. Barahona and J. Christensen (ed). IOS Press. 133-142.
2. Rector, A., W. Solomon, W. Nowlan and T. Rush (1995). A Terminology Server for Medical Language and Medical Information Systems. Methods of Information in Medicine, Vol. 34, 147-157
3. Rector, A., A. Gangemi, E. Galeazzi, A. Glowinski and A. Rossi-Mori (1994). The GALEN CORE Model Schemata for Anatomy: Towards a re-usable application-independent model of medical concepts. Twelfth International Congress of the European Federation for Medical Informatics, MIE-94, Lisbon, Portugal, 229-233.
4. Rector, A. (1995). Coordinating taxonomies: Key to re-usable concept representations. Fifth conference on Artificial Intelligence in Medicine Europe (AIME '95), Pavia, Italy, Springer. 17-28.
5. Rector, A., J.E. Rogers, P. Pole (1996) The GALEN High Level Ontology. Fourteenth International Congress of the European Federation for Medical Informatics, MIE-96, Copenhagen, Denmark
6. Rector, A. and W.A. Nowlan (1993). The GALEN Representation and Integration Language (GRAIL) Kernel, Version 1. The GALEN Consortium for the EC AIM Programme. (Available from Medical Informatics Group, University of Manchester).
7. Gaines BR, Shaw ML and Woodward JB (1993). Modelling as Framework for Knowledge Acquisition Methodologies and Tools. International Journal of Intelligent Systems 8(2): 155-168.
8. Campbell KE, Das AK and Musen MA (1994). A logical foundation for representation of clinical data. JAMIA 1(3): 218-232.
9. Cimino J (1994). Controlled Medical Vocabulary Construction: Methods from the Canon Group. Journal of the American Medical Informatics Association 1(3): 296-197.

10. Evans D (1988). Pragmatically-structured, lexical-semantic knowledge bases for unified medical language systems. Proceedings of the Twelfth Annual Symposium on Computer Applications in Medical Care, Washington DC, IEEE Computer Society Press: 169-173.
11. Huff S and Warner H (1990). A comparison of Meta-1 and HELP terms: implications for clinical data. Fourteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-90), Washington DC, IEEE Computer Society Press: 166-169.
12. CEN PT 251:ENV 1828 Surgical Procedure Modelling