

GALEN's Model of Parts and Wholes: Experience and Comparisons

Jeremy Rogers, Alan Rector

Medical Informatics Group, Department of Computer Science, University of Manchester, UK
And: *OpenGALEN* www.opengalen.org

Part-whole relations play a critical role in the OpenGALEN Common Reference Model. We describe how particular characteristics of the underlying formalism have influenced GALEN's view on partonomy, and in more detail discuss how specific modelling issues have driven development of an extended set of partitive semantic links.

INTRODUCTION

Part-whole relations (or 'partonomies') play a critical role in medical concept representation. This is most obvious in the modelling of anatomy but it also affects the representation of surgical procedures – which usually act on anatomy – as well as many physiological or disease processes. Part-whole relations have long been the subject of extensive study both in the linguistic and knowledge representation communities.

OpenGALEN is an open source continuation of the EC funded GALEN programme, developing a generalised and application independent architecture for medical terminology and associated software.

In this paper we present in detail the mechanisms used to represent part-whole relations in the *OpenGALEN* Common Reference Model of medicine¹. We compare our approach with specific mechanisms used by others in the medical and conceptual modelling communities. (Comparison with formal mereology² and lexical semantics³ are beyond the scope of this paper.)

Relevant Features Of The Grail Description Logic

The GRAIL description logic formalism and the issues in its design are described in detail elsewhere^{4,5}. However, certain features of the formalism are specifically relevant to representing part-whole relations. These are summarised here:

- Semantic links ('attributes') can be declared as transitive, *e.g.* to represent formally that 'parts of parts of a whole' are 'parts of the whole'.
- One semantic link can be declared to be inherited along another ('refinement'), *e.g.* to state that 'a disease involving a structure that is a *part-of* a whole' is a kind of 'disease involving the whole'. (This is approximately equivalent to support for 'right identities' as used by SNOMED-RT⁶).

- Semantic links as well as concepts may be organised into an is-kind-of (subsumption) hierarchy, *e.g.* to say that the link component-of is a kind of the link part-of, with the interpretation that a component-of something is *de facto* also a kind of a part-of that thing.

KINDS OF PARTONOMIC RELATIONSHIP

It has long been recognised that several different subtypes of the *part-of*, or aggregating, relation may be identified. A commonly used list in the Artificial Intelligence and conceptual modelling communities is due to Winston⁷ [made more accessible by Odell⁸]. The list as simplified by Odell is shown below:

component	<i>e.g.</i> handle on a car door
stuff	<i>e.g.</i> flour in bread
portion	<i>e.g.</i> a slice from a loaf of bread
area	<i>e.g.</i> city in a country
member	<i>e.g.</i> ship in a fleet of ships
partner	<i>e.g.</i> Laurel in Laurel & Hardy
piece	<i>e.g.</i> handle when removed from door

GALEN set out to make use of the first five; the sixth is a special case of 'member' in which the objects being joined are instances, not classes. 'Partner' is therefore not relevant to the *OpenGALEN* model since this restricts itself to classes. The seventh relation - piece - was explicitly excluded from GALEN. Medicine is strongly grounded on normative anatomy; the piece/component distinction was thought to be counter-productive in this domain. Congenitally variant and abnormal anatomy is modelled using other techniques and links.

We have, however, since found it necessary to elaborate and extend this partonomic link structure. These extensions are required to account for:

- a) patterns of inheritance along part-whole relations
- b) distinctions which are important to medicine but not accounted for in the above scheme

Notation and naming

In this paper semantic links appear in lower case and italicised. The names for each link type used in this paper are different from those in the current *OpenGALEN* model where, for historical reasons, link knowledge names are potentially confusing. A

mapping from the names used in this paper to the real *Open*GALEN links is presented later (fig. 2).

Component

In general, a *component-of* a *component-of* a whole is a kind of *component-of* the whole. Thus, the *component-of* semantic link is fully transitive. However, difficulties arise when considering the inheritance of other, non-partitive semantic links across the *component-of* link.

Functions of substructures appear to be inheritable along *component-of* links: the functions of the whole includes the sum of the functions of the components. For example, the pancreas is a kind of structure that secretes insulin, even though this function is actually fulfilled by identifiable and discrete components within the pancreas, the Islets of Langerhans. But inheritance of *function-of* across *component-of* is only valid up to a certain (often arbitrary) level of anatomical aggregation: the gastrointestinal system as a whole is not considered to be a kind of structure that secretes insulin, even though the pancreas is a component of it. By contrast, abnormality of gastric secretion (e.g. achlorhydria) is a kind of functional gastrointestinal disorder, but thyroid hypersecretion is not a disorder of neck function.

We capture this variable inheritance by preventing inheritance of *function-of* across *component-of* itself but enabling it for a child link of *component-of*: *func-component-of*. Under this scheme, the Islets of Langerhans are *func-component-of* the pancreas but the pancreas is only a *component-of* the gastrointestinal tract. Conversely, the stomach is a full *func-component-of* the gastrointestinal tract.

Stuff

The top level ontology of the *Open*GALEN model^{9,10} makes a primary distinction between structure and substance. In general, substances should be *stuff-of* structures. This distinction is not, however, without its tensions. In the case of haematemesis, is the blood *stuff-of* vomitus or *portion-of* it?

Odell states that the *stuff-of* relation is a subtype of the *part-of* relation. We have observed that it often also strongly implies a *contained-in* relation, particularly in the context of substance *stuff-of* liquid. (e.g. blood *contained-in* vomitus; sodium *contained-in* plasma). This led us to split *stuff-of* into *made-of* and *mixed-throughout*. *Mixed-throughout* is then a descendent of both *part-of* and *contains* through its direct, hybrid parent: *partitively-contains*. It has further descendent links *dissolved-in* and *suspended-in* to represent the difference between a solution and a suspension or emulsion (water *dissolved-in* alcohol

but water *suspended-in* oil).

Inheritance of *stuff-of* over other partitive semantic links also requires care. For example, the fat tissue within an eyelid can be viewed as being *stuff-of* the eyelid. Within the *Open*GALEN model, however, the eyelid is itself separately given as a *component-of* the eye as an organ and a *surface-of* the face. This presented difficulties when considering the surgical procedure to resect that fatty tissue. When the *stuff-of* relationship is inherited over both the *component-of* and *surface-of* relations, this operation is classified as both a kind of 'eye surgery' and 'facial surgery'. This conflicts with usual convention and suggests that inheritance over both links is inappropriate.

Portion

From an early stage we chose to subdivide the *portion-of* relationship along overtly topological lines. Thus, we distinguish between *surface-of*, *segment-of*, *pouch-of*, *layer-of* and *irregular-piece-of* semantic link types. This was motivated partly by a desire to mirror and reinforce awareness by modellers of the topologically based principles by which the anatomy concept space itself was divided¹⁰.

The subtypes of the *portion-of* semantic link, in combination with the GRAIL sanctioning mechanism, are used to more strongly and explicitly enforce the overall topological schema. For example, the model is constrained such that only structures already known to have a linear shape may be subdivided using the *segment-of* semantic link.

Subdividing *portion-of* has also proved valuable in providing context for controlling inheritance of the *portion-of* semantic link or its subtypes. For example:

inferior surface *surface-of*
articular cartilage *component-of*
knee joint *component-of*
knee region *portion-of*
thigh *component-of* lower extremity

is not a kind of :

structure *surface-of* leg

but it is a kind of :


structure *portion-of* leg.

Similarly, a *segment-of* an artery *component-of* the thigh is not a kind of *segment-of* the leg.

The *pouch-of* relation was conceived specifically for the relation between the vermiform appendix and the caecum, or a diverticulum and the bowel more generally, but is also used for e.g. recesses of joint

cavities. Its primary value is to allow indexing of objects with this particular relationship to a larger whole; this group is often clinically significant.

The *layer-of* relationship differs from *irregular-piece-of* in that it specifically implies a particular plane of section, whereas *irregular-piece-of* specifically does not. *Layer-of*, consequently, requires special consideration. Ordinarily this semantic link can not be relied on to be self transitive (even though Odell states that all partitive relations are transitive⁸): in fig. 1 structure (A) is a *layer-of* structure (B). However, as it is taken at right angles to the original plane of section it would not normally be viewed as a kind of *layer-of* the whole (C):



However, in human anatomy, it is usually the case that all successive layers are concentric or otherwise have parallel planes of section (we have not yet found an exception) and the case where one is orthogonal to the super-layer does not arise. This allows the *layer-of* semantic link to be transitive in the *OpenGALEN* model. However, unlike the other subtypes of *portion-of*, *layer-of* can not be inherited over the *component-of* semantic link: a *layer-of* a component is not a kind of *layer-of* the whole, though it is a *portion-of* the whole.

Member

The *GRAIL* formalism and, consequently, the *OpenGALEN* model support only a weak notion of collection. So far, in our modelling of discrete anatomy, we have not found reason to make use of even that limited functionality. Therefore, although the member-of link exists in the model, we have no extensive experience of its use at this time.

Area and Containment

Our experience in the domain of anatomy confirm Odell's observation of an inconstant relationship between *contains* - true topological inclusion - and *partonomy*. The container may be treated as *part-of* the content (pericardial disease is a kind of heart disease; pleurisy is a kind of lung disease). Alternatively, the reverse may be true (colitis is a kind of abdominal disease) or there may be no partitive relationship between the container and its content (cerebral abscesses are not skull disorders).

Thus containment and partitivity should properly be

modelled independently. However, for authoring economy and to work around limitations within the algorithm to expand *GALEN* intermediate representation into *GRAIL*¹¹, a hybrid relationship is used: *partitively-contains* - similar to Odell's area relation - to represent the case where the content is also considered to be *part-of* the container. This semantic link is given as a descendent of both *contains* and *component-of* such that the link subsumption hierarchy is a graph, not a tree. The impact of this convenience, and of a semantic link polyhierarchy more generally, on *GRAIL*'s compute performance is under review.

Connectivity

Connective relationships are discussed in Odell's presentation of partitive relations, where he notes that *connects* does not guarantee *part-of*. Odell suggests that a connector can only be *part-of* the connectee if it also provides functional support for it.

In *GALEN*, *connects* has proved useful for describing e.g. the relationship between the carina and the trachea and main bronchi or between the major heart valves and the chambers of the heart and/or major vessels they are situated between.

Our observations confirm an inconstant relationship between *connects* and *part-of*. For example, a saddle embolism at the aortic bifurcation is a kind of disorder of the abdominal aorta (connector is treated as *part-of* the connectee). By contrast, ulceration at a colostomy is not a kind of colonic ulceration (the connector is not *part-of* the connectee).

A particularly complex case was encountered in modelling the cardiac septum. This single structure *connects* the left and right heart but it is *part of* neither, being a *component-of* the heart as a whole rather than of one or other half. However, it also has discrete *components-of* itself (the bundle branches) that are considered to be *part-of* one or other of the left or right ventricles respectively.

Again, similar to our treatment of containment, we presently use a hybrid semantic link *partitively-connects* for economical authoring of the case where the connector is also *part-of* the connectee.

Arbitrary part

In our representations of surgical procedures and of diseases we encountered the notion of apparently discrete, well circumscribed - and well understood - structures whose parts are not related in the normal anatomical sense. For example, hand-foot-and-mouth disease (caused by Coxsackie A-16) is typically characterised by a rash that affects the hands, feet and mouth while sparing the structures in between. Several representational choices present themselves

in this situation: is this one rash with multiple locations or three rashes with one location each?

We also experimentally considered one rash with one location, but this required a mechanism to represent the notion of a single, anatomical structure with component parts the hands, the feet and the mouth. We felt the *component-of* relation did not seem appropriate for this purpose and so a dedicated link – *arbitrary-part-of* – was created specifically for the purpose. Such a structure would be defined as:

```
ArbitraryBodyConstruct
  which arbitrary-part-of hand
  which arbitrary-part-of foot
  which arbitrary-part-of mouth
```

..where ArbitraryBodyConstruct is the parent of all such structures. Similar structures might be ‘head and neck’, ‘hand and wrist’ or ‘retina and choroid’ but, although some of these are already present in the OpenGALEN model, their representations do not yet consistently use the ArbitraryBodyConstruct method.

A variant of this problem occurs with adjacent structures. We speak of “A fracture of the tibia and fibula” but would not talk of a single fracture of the “femur and humerus”. We normally expect a single fracture to have a single location, but for adjacent structures we allow it to have two. This contrasts with the rash example, where the locations of the rash are not required to be adjacent.

Additional abstractions

In addition to deepening the original Winston-Odell structure, primarily by adding sub-types of *portion-of*, we have also added a number of additional abstractions in anticipation of retrieval and aggregation requirements.

Our initial division of part-of is into *structure-part-of* (which subsumes the flavours of *portion-of* and *component-of*) and *stuff-of*. Additionally, *layer-of* and *irregular-piece-of* share an added common parent that represents their disjunction: *solid-piece-of*.

Subprocess

Winston-Odell’s original structure deals with part-whole relationships between structures and substances. OpenGALEN’s full partonomic model also recognises a *subprocess-of* relationship between processes and subprocesses e.g. between the first stage of labour and parturition as a whole, or between abdominal straining and coughing.

We have also experimented with the notion of phases of a process and with episodes of processes.

THE OPENGALLEN PARTONOMIC LINKS

The discussion above presents the case for an ontology of partitive semantic links that is broader and deeper than most. The knowledge names used throughout are different from those actually used in the *OpenGALEN* model. Figure 2 shows the complete *OpenGALEN* ontology of partonomic semantic links as an expanded tree structure, together with the mapping of each such link to the names used in this paper.

Note that an additional level of the tree, present in the full *OpenGALEN* model, is not shown here. This final layer – members of which are denoted by the infixes ‘specific’, ‘specifically’ or ‘unique’ - is an artefact of the limited implementation in GRAIL of link number restriction and is not unique to partonomic modelling¹⁰.

Figure 2

InversePartitiveAttribute	<i>part-of</i>
InverseStructuralPartitiveAttribute	
isDivisionOf	<i>structure-part-of</i>
isArbitraryComponentOf	<i>arbitrary-part-of</i>
isLinearDivisionOf	<i>segment-of</i>
isSolidRegionOf	<i>solid-piece-of</i>
isBlindPouchDivisionOf	<i>pouch-of</i>
isLayerOf	<i>layer-of</i>
isSolidDivisionOf	<i>irregular-piece-of</i>
isStructuralComponentOf	<i>component-of</i>
isFunctionComponentOf	<i>func-component-of</i>
isPartitiveConnectionOf	
isPartitivelyContainedIn	
isMixedThroughout	<i>portion-of</i>
isDissolvedWithin	<i>dissolved-in</i>
isInSuspensionWithin	<i>suspended-in</i>
isSurfaceDivisionOf	<i>surface-of</i>
makesUp	<i>stuff-of</i>
InverseProcessPartitiveAttribute	
isSubprocessOf	<i>subprocess-of</i>

OTHER WORK

The Digital Anatomist Foundation Model¹² and the *OpenGALEN* model share similar high level distinctions in their respective concept spaces; substances are separated from structures, and structures are further divided between space and physical entity. Some lower level distinctions are also similar, for example anatomical junctional structures are identified as a class. The Digital Anatomist semantic link ontology includes *component-of*, *stuff-of*, *layer-of* (lined-by), *contains* and *connects*, although there appears to be some overlap or redundancy in the link set: ‘consists of’ and ‘is divisible into’, for example, appear to be equivalent.

Padgham¹³ noted link inheritance combinations which follow the pattern “the hand is part of the upper extremity, therefore the skin of the hand is *part-of* the skin of the upper extremity”. The best GALEN’s rules can do is that the skin of the hand is a *kind-of* the skin of the upper extremity. In practice the inheritance patterns are sufficiently similar that no significant errors in classification have been found within the applications so far encountered.

Udo Hahn^{14,15} has formalised a commonly used alternative mechanism for inheriting links over one another: for each whole, a corresponding concept ‘part-of-that-whole’ exists. A detailed comparison has not been made, but this approach appears to achieve many of the same aims while making fewer special demands on the description logic.

DISCUSSION

We have shown how an initially simple set of partonomic semantic links has been extended during the development of the *OpenGALEN* model. The primary driver has been coping with subtleties observed in link inheritance, particularly when partitive semantic links form part of a chain of links across which various links are then inherited. Additional influences have been the need to index clinically significant classes by particular subtypes of partitive relationship.

It might be argued that an alternative mechanism could allow the same problems to be addressed while allowing a simpler ontology of partonomic links. For example, link inheritance rules might be constrained by specific combinations of links and the concept classes they link, rather than only of links (as now).

However, each of our partitive semantic links is motivated by specific exceptions to inheritance rules and reflects real features of our conceptualisation of anatomy. We believe a simpler link ontology would not obviate the need for any apparent complexity but instead require it to be moved elsewhere.

The *OpenGALEN* model of anatomy is highly multiaxial, but has been constructed using very few asserted multiple parent relationships: almost all multiple parent (is-a) relationships are inferred from other explicit, and hence reusable, information present separately in the model.

The classificatory inference process makes extensive use of link inheritance, often over long chains of successive links. Fine control of link inheritance, implemented in part through a rich partitive link ontology, allows this classificatory reasoning to be performed safely.

Acknowledgments

With thanks to all in the GALEN-IN-USE consortium.

GALEN-IN-USE is funded as part of Framework IV of the EC Healthcare Telematics research program

References

1. <http://www.opengalen.org>
2. Sowa, J. (1999). Knowledge Representation, Moorage
3. Cruse, D. A. (1986). Lexical Semantics. Cambridge, New York
4. Rector, A., Bechhofer, S., Goble, C., Horrocks, I., Nowlan, W. and Solomon, W. (1997). The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine* 9;: 139-171.
5. Rector, A.L., Zanstra, P. E., Solomon W.D., Rogers J.E., Baud, R., et al. (1999). Reconciling Users' Needs and Formal Requirements: Issues in developing a Re-Usable Ontology for Medicine. *IEEE Transactions on Information Technology in BioMedicine* 2(4);:229-242.
6. Personal communication
7. Winston, M., Chaffin, R., and Herrmann, D. 1987. A Taxonomy of Part-Whole Relations. *Cognitive Science*, 11;: 417-444.
8. Odell, JJ. (1998) Six different kinds of aggregation. In *Advanced Object-Oriented Analysis & Design Using UML*. Cambridge University Press;:139-149
9. Rector, A.L., Rogers J.E., Pole P.M. (1996) The GALEN High Level Ontology. J. Brender, J.P. Christensen, J-R. Scherrer and P. McNair (Eds.) *Medical Informatics Europe '96 (Part A)*, IOS Press;:174-178
10. Rogers J.E., Rector, A.L. (1997) Extended Core model for representation of the Common Reference Model for procedures. <http://www.opengalen.org/open/crm/CRM-documentation.zip>
11. Rogers J.E., WD Solomon, Rector, A.L., Pole P.M., P Zanstra, E van der Haring, (1997) Rubrics to Dissections to GRAIL to Classifications C. Pappas et al. (Eds.) *Medical Informatics Europe '97*, IOS Press;:241-245.
12. Rosse, C., Shapiro, L., & Brinkley, J. (1998). The digital anatomist foundational model: Principles for defining and structuring its concept domain. *Annual Fall Symposium of American Medical Informatics Association, Orlando FLA*. Hanley & Belfus Inc. Philadelphia PA;:820-824
13. Padgham, L. and Lambrix, P. (1994). A Framework for Part-of Hierarchies in Terminological Logics. KR-94
14. Hahn, U., Schulz, S. and Romacker, M. (1999). Partonomic reasoning as taxonomic reasoning in medicine. *Proc. of the 16th National Conf. on Artificial Intelligence & 11th Innovative Applications of Artificial Intelligence (AAAI-99/IAAI-99)*, Orlando FL, AAAI Press/MIT Press.
15. Hahn, U., Schulz, S. and Romacker, M. (1999). Part-whole reasoning: a case study in medical ontology engineering. *EEE Intelligent Systems & their applications* 14(5);: 59-67.