

Metonymic and Holonymic Roles and Emergent Properties in the SNOMED CT Ontology

Jon Patrick

Centre for Health Informatics Research & Development

School of Information Technologies

University of Sydney

Sydney Australia

{jonpat@it.usyd.edu.au}

Abstract

This paper discusses the manner in which SNOMED CT (SCT) has confused the metonymic role of some class labels as holonyms and has inappropriately assigned property inheritance down a holonymic chain due to its transitivity. The notion of emergent properties is introduced as the only form of property that can exist on a holonym and its use in a hypernymic inheritance hierarchy is discussed. The consequences of this modelling approach for SCT are discussed and the use of metonymic substitution for holonyms at the point of care is presented as a source of confusion for causing the modelling of inheritance in holonymic hierarchies for clinical care. The mathematical modelling of the metonymic substitution is discussed but left as future work.¹

1. Introduction

SNOMED CT (SCT) is a very large scale ontology used for the description of certain classes of medical and health community knowledge. We are interested in its value for the functional purposes of computation at the point of care and the issues in maintaining and delivering it for those purposes. SCT is maintained by the College of American Pathologists (CAP) and not necessarily fit for purpose for all health and medical domains. Although very wide ranging with over 360,000 concepts and 1.2 million relations it lacks some fundamental facets of a full ontology. It is not an ontology in the strictest sense of the word as it contains a great deal of knowledge that is compositions of fundamental elements within the system. That is, it is possible to express the same collection of ideas with multiple strategies either as a single concept (if sanctioned) or as a group of individual concepts. Also, classes do not have

attributes but rather they are expressed through relations between “characteristics” and “concepts”. SCT although maintained using the Protégé logic engine is released as a set of 3 files prepared as comma separated values (CSV) and so fails to come with a logic engine to preserve its logical structure and constraints. Lastly the modelling has been performed over more than 40 years and in the late 1990s it was merged with another large ontology the Reed codes developed in the UK. This has blended together two different modelling paradigms so that SCT has a significant admixture of concepts and relations which have lead to redundancy, rival explanatory theories, and confused modelling.

Ontologies to be useful for computation for practical tasks need to be precise in two ways; firstly, they need to match as closely as possible our understanding of the natural world we wish to deal in. Ontologies need to be accurate and closely follow human understanding otherwise they create confusion in their design and their use. This means the ontology needs to be constructed with very close attention to the meanings of the terms used in it. Hence the names used in the ontology need to represent closely the understanding we have of the real world.

Secondly the variety of linguistic usage of those terms needs to be explored to uncover the diversity of semantic roles of the terms and ensure that only those roles that are useful are included in the ontological modelling, and to remove ambiguity in the use of those roles. Precision of definition is particularly important in establishing the relationships between elements and identifying the fundamental atomic elements and how they combine together systematically to make more complex semantic concepts.

Ontologies also consist of abstractions with each level of the ontology being more abstract as one moves up the hierarchy. This structure allows us to talk about the world at the different levels of abstraction. In using an ontology for computation there are two basic forms of abstraction available, aggregation and generalisation. Generalisation hierarchies are used throughout SCT as the basic mechanisms for relating content. Aggregation hierarchies on the other hand have not been used properly but rather transposed so that they appear

¹ Copyright (c) 2006, Australian Computer Society, Inc. This paper appeared at *The Australasian Ontology Workshop (AOW 2006) Hobart, Tasmania, Australia*. Conferences in Research and Practice in Information Technology, Vol. 72. Mehmet Orgun & Thomas Meyer, Eds. Reproduction for academic, not-for profit purposes permitted provided this text is included.

like generalisation hierarchies. We investigate this replacement for a proper aggregation hierarchy and argue that it comes from a misunderstanding of both the linguistic use of terminology at the point of clinical care and the logic arguments developed for its justification. In particular in SCT the holonymic (or hypernym) role of an aggregating concept is used as a source of inheritance which is clearly incorrect. Our explanation for this SCT modelling strategy is that the role of such a holonym has undergone the process of metonymic substitution, which is substitution of the authentic word for one that serves as a metaphor for the original. The assignment of attributes and relations of the meronymic (sub-part) members of the holonym (super-part) to be one of the holonym itself can at best be called metonymic inheritance. Importantly, if it is allowed to operate at all, it must operate from the bottom up, that is the attributes move from the part to the whole, that is, in reverse to what we normally think of as the direction of inheritance, from the top down.

2. The Meaning of Aggregation and Generalisation.

The discussion of Aggregation and Generalisation is drawn from the database literature so as to emphasize their computational aspects. Smith & Smith make the basic statement “aggregation and generalization are independent activities” (1977, p116)¹. Although they are two fundamental means of organizing the knowledge about the relationships between concepts or entities, they specify knowledge that is independent of each other. Generalisation is a hierarchy of classes and subclasses in a subtype relationship, usually referenced as IS_A. Aggregation is a hierarchy of components sometimes called a PART_OF or PART-WHOLE relationship. Generalisation and Aggregation are *relationship types* And hence cover many instances of different hierarchies of these *types* found in natural language. The similarities and differences between these two types of relationships are fundamental to understanding how ontologies are organised and the limitations to their processing.

Generalisation is a hierarchy of classes and subclasses where each subclass has a transitive relationship with its parents and children, that is, properties are carried from subclass to subclass down the hierarchy, but also the properties introduced into a subclass that are separate to the inherited properties are *not inherited* back up the hierarchy.

Aggregation is the idea that individual types of parts are brought together into a hierarchy to make a whole, such as in the assembly of a motor car. Hence, a sparkplug is part of a motor, is part of a car. Importantly there is a transitive relationship between the parts, that is, the sparkplug is part of the car. However, there are no inherited properties, so

red wheels do not make the car a red car, nor does a red car make the wheels red. Furthermore, the relation is *irreflexive*, that is, an entity cannot have a part-of relationship with itself, so a car is not part of a car.

Algebraically the transitive property can be expressed as:

If aRb, bRc, then aRc, where R is a defined relationship such as IS_A_SUBCLASS_OF, or IS_PART_OF.

In this example the relation IS_A_SUBCLASS_OF is defined to be of the relationship type Generalisation, while the relation IS_PART_OF is relationship type Aggregation. So whilst both relations are transitive and we can apply that algebraic function to their hierarchies, the meaning content of the relationship names are different and the relationship types are different.

3. Linguistic Terminology

In linguistics the Generalisation relationship is known as a *hypernym-hyponym* relationship. Aggregation is known as the *holonym-meronym* relationship. A further linguistic concept, *metonym*, is needed to complete the explanation of the problems with the SCT structures. A *metonym* is a word used in place of the correct word because it is associated with it in some way, so the word “dish” can be used in place of “roast beef” when talking about the evening meal. This is a very common linguistic strategy and in one of its extreme forms it is called a *euphemism*.

4. Processing with Ontologies

In SCT there are 19 generalisation hierarchies such as Organism, Body Structure, Clinical Finding, etc. The concept “part-of” is defined in the Attribute generalisation hierarchy. It has a means of representing part-whole concepts by using SEP structures based on the ideas created in Schulz, Romaker & Hahn² which in turn claims its heritage in the GALEN project Rector, Bechofer, Goble, Horrocks, Nowlan & Solomon³. These papers present the problem of part-whole reasoning as a problem of differentiating it from reasoning for subsumption. They point out that some systems, MESH for example, treat reasoning for generalisation relations and part-whole relations the same way producing inconsistencies such as *blood* being a hypernym for both *foetal blood* (a type of blood) and *blood plasma* (a part of blood). They point out that a different logic is needed to compute over a part-whole hierarchy which consists of the features *transitivity* and *part-whole specialisation*.

5. Transitivity

A key issue for Schulz et al is the notion of whether the meronymic relationship (IS_PART_OF)

is transitive as in the relationship IS_SUBCLASS_OF (or IS_A). The example provided is
 appendix IS_PART_OF colon
 colon IS_PART_OF intestine
 as this is manifestly true for common sense knowledge then it must be true that
 appendix IS_PART_OF intestine
 hence transitivity is demonstrated.

However Schulz et al do not deal with the issue of inheritance of properties in either a modelling or a linguistic sense for a Generalisation hierarchy. Importantly later in their paper they assume that given the transitivity of both relationship types (that is generalisation and aggregation) that any property is also inherited between types. This is manifestly untrue. Take the example of a motor car and its wheels. A blue motor car does not necessarily have blue wheels and a car with a set of racing wheels is not necessarily a racing car. Hence the characteristics of each part are not necessarily inherited either up or down a holonymic hierarchy, in fact that is the point of differentiation between the two types of relationships, hyponyms have inheritance of characteristics and meronyms don't have inheritance. Furthermore red upholstery in a red car are not instances of the same red characteristic, they are merely coincidental facts about the car. While it is possible to assign algebraic symbols for all these phenomena and then to perform computational manipulation that doesn't mean that any computation is sensible. Moderation of what is sensible to compute must lie with the meaning intention of the symbolic representation.

The misconception of the sharing of inheritance due to the equivalence of transitivity between the two relationship types needs to be investigated both from a linguistic perspective and from an historical perspective to understand how their interpretation arose.

6. The Usefulness of the "Part-Whole Specialisation"

Schulz et al take holonymic *specialisation* from Horrocks et al⁴ and call it *part-whole specialisation*. This phrase creates a warning sign of future difficulties. Specialisation is the inverse of generalisation, and generalisation-specialisation are distinctly different to part-whole relationship, so the term *part-whole specialisation* appears to be a contradiction in terms.

Schulz et al produce a holonymic example with "shaft of femur" is PART_OF "femur" (Fig 1), in a parallel structure they produce the pair "fracture of shaft of femur" IS_A "fracture of femur" as a specialisation. They then assert that not only are the structures related by the relationship "FRACTURE_OF" (relations R1 and R3) but the hyponym "fracture of shaft of femur" is related to the holonym "femur" by the same relation, R2, that

is "fracture of shaft of femur" is a FRACTURE_OF "femur".

There are a number of difficulties with the configuration in Fig 1 and the line of argument presented in Schulz et al. The logical deduction follows these lines:

Premiss: the "femur" and "shaft of femur" have a holonym-meronym relationship,

Deduction:

If the hyponym to meronym have a relationship, in this case FRACTURE_OF,

then the hyponym to holonym have the same relationship (R2),

and thereby the meronym is a *specialisation* of the holonym.

furthermore by asserting the premiss that the hypernym has the same relationship to the holonym(R3) as the hyponym has with the meronym (R1), that a hypernym-hyponym relationship is thereby verified.

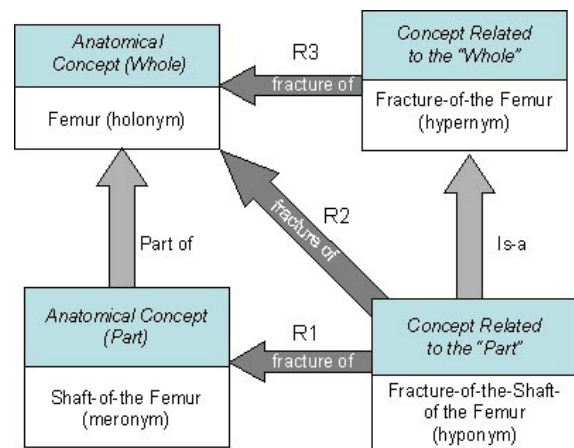


Fig 1. Part-Whole Specialisation as represented by Schulz et al. R1, R2, R3 are the relationship "fracture of". R2 is a relationship from the hyponym to the holonym.

A general form of this argument is embodied in the SEP structure as used to model an ANATOMICAL-PART-OF relation for the relationship between physical parts of an organism (*op cit*). It consists of three nodes: S, E and P nodes. S is a structure node which subsumes the E-node and the P-node and there is an ANATOMICAL-PART-OF relation from the P-node to the E-node (fig2).




Fig 2. Structure of S-node for modelling Part-Whole Specialisation (after Schulz et al).

The E-node denotes the *whole* anatomical part being modelled and the P-node denotes any part of the E-node entity. The P-node is then the parent of the same structure repeated down a generalisation hierarchy, so:

the intestine-structure (S-node) has an intestine (E-node) and intestine-part (P-node) which has a colon-structure (S-node) which has a colon (E-node) and colon-part (P-node) which has a caecum-structure (S-node) which has a caecum (E-node) and caecum-part (P-node) which has an appendix-structure (S-node) which has an appendix (E-node).

Schulz et al justify this structure in saying “by introducing a special data structure for part-whole encoding we build up specifically structured IS-A hierarchies which support the emulation of inferences typical of transitive PART-OF relationships. The same formalism allows for conditioned part-whole specialization”. They have in effect created a data structure that masks the PART-OF relationship between anatomical elements on the argument that they don’t have a transitive relationship (whereas they are transitive) so as to allow for inheritance (which is not permissible).

However they concede that the inheritance of the specialisation relationship down the holonymic hierarchy is not always found to be consistent with medical knowledge. The cited contrary example is:

While

Perforation of Appendix IS_A_SUBCLASS_OF Intestinal Perforation **is true,**

it is not true that,

Inflammation of Appendix (Appendicitis) IS_A_SUBCLASS_OF Intestinal inflammation (Enteritis).

Their arguments for introducing specialisation into the Aggregation hierarchy are not convincing and do not account for a number of matters and introduce the possibility of significant errors of reasoning. It appears that two other motivations for this configuration exist (personal communication):

that “part-of” is meant to provide for referencing a portion of a body part that has no individual label, for example where the portion of the lung (say 35%) is removed, and also,

the logic engine used to verify the hierarchies only provided for transitivity verification through the implication operator in the generalisation hierarchy and not through the aggregation hierarchy.

Both of these explanations provide some reasoning for the structure and also create a motivation for wanting to review the solution to get a representation closer to a satisfactory ontological form. Considering the problem from another

perspective we can say, firstly, a “fracture to the shaft of femur” is not a type of “shaft of a femur”, that is, it is not a subclass of either a “femur” or “shaft of femur”. To make an analogy, a broken car is not a broken wheel and a broken wheel is only a broken car through a metonymic substitution where the correct word, “wheel” is substituted by a word associated with it, “car” in this case. If the “shaft of the femur” is broken then it is not the “femur” that is broken but rather the “shaft” and on the basis of non-inheritance between holonym and meronym the one is not the other. The key confusion occurs because the term “femur” which is a holonym in the hierarchy is used as well as a metonym, that is, as a substitute for the correct word by virtue of association with it.

One weakness in the argument lies in the nature of the modelling of the physical world. A fracture to anything is not an entity of itself but a CHANGE OF STATE of an entity and is better modelled that way. No doubt clinicians at the point of care speak about “a fracture to the shaft of the femur” without paying attention to the linguistic structure. The “fracture to” is really serving as a descriptor of the entity as in “a fractured shaft of the femur” where the participle “fractured” is clearly serving in the epithet role.

Turning to the logic of the argument for the part-whole specialisation, there is no credible basis for any steps in the deduction. The argument hinges on the notion that “specialisation” exists because a hypernym has the same relationship to each of the members of the holonym-meronym pair. However the argument relies on deducing that the relationship between hyponym and holonym is created by virtue of the relationship of hyponym to meronym, then meronym to holonym –this is a specious argument.

Our analysis questions the applicability of the SEP model to a holonymic structure and demonstrates that it is not applicable at least in the scope of the examples used so far. The explanation for the justification used by Schulz et al is shown to be flawed because they shift from defining a term in a holonymic role and then change the role in usage of the holonym to a metonym.

Schulz et al use another example to discuss the generalisation of their SEP model where their process is also manifest. They assert that if INTESTINE-STRUCTURE is an S-node then a PERFORATION-OF-APPENDIX is also a PERFORATION-OF-INTESTINE as Appendix IS-A-PART-OF Intestine. This is the same faulty reasoning. It is true that if the holonymic tree is only defined as far as “intestine” then there is no such thing as a PERFORATION-OF-APPENDIX and any perforation of anywhere in the Intestine is a PERFORATION-OF-INTESTINE. However if the holonymic tree is further described in all of its parts then the word “intestine” can only be used in a metonymic role and as such does not have any of the characteristics of the real perforated part, that is, it is

acting as a surrogate. Taking the example of a car-motor-spark_plug meronymy, if the spark plug is cracked then the motor won't work and the car won't work, but neither the motor nor the car is cracked. Both the motor and the car have emergent properties, that is properties they have by virtue of being a "whole" that none of their parts have, e.g. a motor runs and a car runs or works or can be driven (in some contexts). In English we have words that differentiate between a cracked spark_plug and a motor that doesn't work and car that's broken down that help create the separation of metonymy from meronymy - we would not (normally) say "My car is cracked" or "My motor is cracked".

This argument does not discount the legitimate use of the PERFORATION-OF-INTESTINE when the meronymic component is unknown as might well happen at point of care. But that situation is as above where the word is being used in the role of a metonym and not as a holonym. This issue leads into the further example used by Schulz et al where they have at the top of a diagram an example of the S-node Intestine structure which descends down an IS_A hierarchy successively to colon, caecum, and appendix. They assert that "enteritis" has the relation INFLAMMATION-OF with "intestine", and "appendicitis" has the relation INFLAMMATION-OF with "appendix" yet "appendicitis" cannot be classified as a type of "enteritis", even though the "appendix" is a PART_OF the "intestine". Importantly they argue that each of these are attached to E-nodes and thereby do not automatically have a subsumption relationship as is true in real life, whereas they claim the Perforation example is a true subsumption and is attached to the P-nodes.

In such a structure if one were to assign to the intestines a property of 8metres long, as might well happen in an operational information system, we would have a colon, caecum and appendix all 8 metres and a very crowded pelvic zone. In terms of diseases in this diagram Enteritis is attached to the Intestine and Appendicitis to the Appendix, both attached by the relationship INFLAMMATION-OF. Apart from the convention that enteritis is a condition of the small intestine and colitis of the large intestine, the naming convention does not match any sensible characterisation of the real world. Furthermore the weakness in the SEP model is exposed. If the correct level of granularity is a match of the disease instance level to the anatomical location, that is Appendicitis IS_AN_INFLAMMATION_OF Appendix, then Appendicitis is a specialisation of the class INFLAMMATIONS_OF_THE

_INTESTINAL_TRACT, that is the correct level of generalisation of the terminology that must match the correct level of meronymy, but nevertheless importantly they are independent of each other. So to say a patient has "intestinal inflammation" is

holonomically correct, but to say they have an "inflamed small intestine" is a metonymic use of "inflamed", the meronymically correct expression being "enteritis of the small intestine". Hence on Schulz et al's diagram the correct entry for Enteritis should be Intestinal_inflammation a sub-class of inflammations that has members {enteritis, colitis, appendicitis} each of which are DISEASES_OF the anatomical components of the Intestine (small intestine, large intestine and appendix) respectively.

7. A Confluence of Aggregation and Generalisation of Hierarchies

A useful example to analyse the confluence of holonymy and metonymy hierarchies is the case of the "intestine" where each part supposedly inherits the characteristic of "hollow structure".

The principle parts of the intestines are the small intestine (duodenum, jejunum and ileum), and the large intestine (appendix, caecum, colon and rectum). At the common sense level each of these structures is indeed "hollow", but more relevantly open at both ends and so rightly can be assigned that characteristic, although the appendix is open only at one end. However common sense should be the last refuge for constructing descriptions in building an ontology. Importantly the intestine components are contiguous, that is they abut one another and they are open at each end and so form a continuum, which might lead one to presume that the intestines really do have the attributes of a hollow structure and open at both ends, that each component inherits. However a closer inspection shows a different story. We can say that the appendix does have a hollow structure, but is not open at both ends and so would that then negate the assertion that the intestines have a hollow structure open at both ends? The answer is clearly no, as the intestines would continue to perform as expected allowing food to pass through. So in fact the components do not inherit the characteristic from their holonym superior. Rather the "inheritance" is in the other direction and is not a true inheritance at all. The intestine has a hollow structure by virtue of its components all having the same structure (hollow and open) AND because they exist in a contiguous configuration. The meaning of the expression "the intestine has hollow structure" is a shorthand for the saying each of the components have a hollow structure open at both ends and they are contiguous, and is in fact a metonymic use of the concept "intestine" not a holonymic use.

A more transparent example can be found from common experience. If I say "I have a red car" then the most common interpretation is the paint colour of the car exterior is mostly red. It is most unlikely that a listener would interpret that the wheels are red. On the other hand if I say "my car is completely red" then the listener is likely to assume the car body and interior decoration are red, with greater

uncertainty as to whether the wheels are red or not. The listener would certainly not expect the engine or the underbody to be painted red. So in the car example the components are red by virtue of a truth statement about each component individually, but the red of each component is not inherited from the holonym (whole car), but rather they are each individually red. This interpretation is more easily recognised with the car example as there is no ready metonymic equivalent usage available in English as there is in the case of “intestines”. Hence it is much more difficult to construct the fallacy of inheritance by holonym-meronym transitivity.

Although ontological modelling does not allow inheritance in a formal sense we can define the concept of *metonymic inheritance*, which is inheritance by virtue of metonymic substitution of the whole for the parts, and is thereby a linguistic usage phenomena rather than an intrinsic ontological construct. In the intestines example this form of inheritance is from all of the parts to the holonym but we cannot rule out the possibility of metonymic inheritance from a single meronym rather than all parts of the holonym. It remains to be determined when such a construct would be useful and what sort of an algebra might be suitable for describing operations on it.

The modelling consequences for this metonymic interpretation is that devices such as SEP are not useful in modelling as they create redundant features and distort the picture as to what is identical but not inherited, that is replicated, from what is identical because it is inherited, that is duplicated.

The argument presented herein leaves an open question as to what might be an intrinsic property of a holonym given that they cannot be inherited. The only possibility can be an *emergent property*, that is a property that exists because of the intrinsic wholeness of the holonym and not a property of any of its meronyms. Such properties are common, for example, the human body is mobile or can move under its own power as we understand with a motor car. Emergent properties cannot be inherited by their nature and hence cannot be inherited down an aggregation hierarchy but may well be inherited throughout a generalisation hierarchy. They exist because there is some aspect of the whole which is greater than the sum of the parts and so are attached to holonyms. Hence the confluence of aggregation and generalisation hierarchies occurs at the point of a holonym which has emergent properties. The emergent properties carry down the inheritance structure of the generalisation hierarchy but not down the aggregation hierarchy.

It remains a task for this model to be used to frame an algebra so as to provide a formal computational mechanism for applying metonymic substitution for logical deduction along the lines of Padgham & Lambrix's⁵ work for holonym-meronym relationships. It is not immediately

evident how such an algebra would function but it must provide for any rank shift in holonymy to trigger the same in any attached relationships, and for emergent properties to be definable but inheritable only in a generalisation hierarchy. Under such circumstances it is also likely that a rank shift of the attributes has to be made for accurate modelling, that is movement from Appendix to Intestinal Tract requires the descriptor of disease to make a likewise rank shift up an aggregative terminology, for example from Appendicitis to Inflammation. In this manner we avoid saying Appendicitis is an inflammation of the intestine, but rather aggregate all the types of inflammation of intestinal parts under a single rubric “Inflammation of the Intestine” the members of which are {Appendicitis, Colitis, Enteritis...}. Then their use as a descriptor of a disease of the intestine can be identified as a metonymic use and the deduction that the true referent is a meronym, perhaps some distance down an aggregation hierarchy, can be computed reliably, and lead subsequently to semantically sensible responses and prompts from an operational decision support system.

8. Meronymy and Metonymy at Point of Care

The difference between using words in a holonymic role and metonymic is important and context dependent. If a patient attends a clinic and reports “I have a broken leg” the clinician will say “we need to X-ray that and find out what is broken”. On the return of the X-ray the clinician will say to the patient “the femur is broken” however to an attending surgeon they will say “the shaft of the femur is broken” this being an important distinction to the “neck of femur being broken” as they have two entirely different treatment requirements. Once that definition is established they can readily revert to talking about the “fractured femur” without any risk of misunderstanding because they know they are really talking about the “shaft of the femur”. The expression “femur” is thereby being used as a metonym and not a holonym.

The question remains as to the level of granularity a meronymic hierarchy needs to be referenced in the point of care. The answer is resolved clearly by the detail at which the care regime is determined. If the difference in care is defined by differentiating between “shaft of femur” and “neck of femur” then the clinician(s) need to establish their reference terminology to that level, and use of the terminology at the level is literal and meronymic, however references to terms above that meronymic level are not holonymic but metonymic, unless otherwise asserted within the context, and hence the metaphorical role of the metonym needs to be accounted for.

None of these arguments are meant to deny the need for clinicians to talk at various levels of abstraction for different listener communities.

Rather the arguments are intended to direct the need for the ontology to catch all aspects of the language variations in systematic ways that exploit not only our knowledge of medicine but our knowledge of language and our tools of logic (and sometimes statistics).

9. Conclusions

We have argued that the SEP model creates contradictions to common sense knowledge. What has been provided as a computational function in Cyc has been taken as a generally applicable principle in GALEN and then applied without consideration of its semantic constraints.

A consequence of this analysis is that the SEP model is not needed to model the transitivity relationship in a holonymic-meronymic hierarchy as it is intrinsically transitive and such usage sets a dangerous precedent of allowing attribute inheritance in a holonymic-meronymic hierarchy which can readily lead to nonsensical assertions.

The identification of the appropriate level in the holonymic-meronymic hierarchy for clinical reference needs to be at the point at which clinical care is determined, and references to holonyms at that point are literal. References about that point are to be considered metonymic unless otherwise resolved.

The confluence of generalisation and aggregation hierarchies occurs at the point at which a holonym has an identifiable emergent property that is inheritable by its sub class members in its generalisation hierarchy. Any application of properties of meronyms to holonyms can only be in the use of the holonym name in a metonymic role and hence the property is not a true property of the holonym but rather a pseudo property or a “metonymic property”.

Knowledge Representation Meets Databases, pp 24-28, 1996.

⁵ Padgham, L. & Lambrix, P. A framework for part-of hierarchies in terminological logics. J. Doyle, E. Sandewall and P. Tarasso (Eds.). Principles of Knowledge Representation and Reasoning. Proceedings of the 4th international Conference (KR'4), Bonn, Germany, 1994.

¹ Smith, J.M. & Smith, D.C.P. Database Abstractions: Aggregation and Generalisation, ACM Transactions on Database Systems, Vol.2 , No. 2, pp105-133, 1977.

² Schulz, S., Romacker, M. & Hahn, U. Part-Whole Reasoning in Medical Ontologies Revisited – Introducing SEP Triplets into Classification-Based Description Logics. Proc AMIA. 1998.

³ Rector, A.L., Bechofer, S., Goble, C.a., Horrocks, I., Nowlan, W.A. & Solomon, W.D. The GRAIL concept modelling language for medical terminology. *Artificial Intelligence in Medicine*. 9:139-171, 1997.

Rogers, J. & Rector, A. GALEN's model of parts and Wholes: Experience and Comparisons. In Overhage, J.M. (Ed.) Proc American Medical Informatics Assoc Symp. 714-8, 2000.

⁴ Horrocks, I., Rector, A. & Goble, C. A description logic based schema for the classification of medical data. KRDB'96 –Proc of the 3rd Workshop on