

Relations in Anatomical Ontologies*

Fabian Neuhaus Barry Smith

1 Introduction

It is now increasingly accepted that many existing biological and medical ontologies can be improved by adopting tools and methods that bring a greater degree of logical and ontological rigor. In this chapter we will focus on the merits of a logically sound approach to ontologies from a methodological point of view. As we shall see, one crucial feature of a logically sound approach is that we have clear and functional definitions of the relational expressions such as ‘*is_a*’ and ‘*part_of*’. While this chapter is mainly concerned with the general issues of methodology, the chapter of this book on ‘Spatial Representation and Reasoning’ [1], will apply the methodology to the specific case of spatial relations. Although both chapters are self-contained, we recommend that they be seen as forming a unity.

2 The semantic content of type terms

The reason why logical rigor is crucial for the development and use of biomedical ontologies becomes clear if we consider their purpose and mode of operation. The term ‘ontology’ is used very ambiguously, but in the life sciences ‘ontology’ means roughly: ‘controlled vocabulary in computer interpretable form’ and in this chapter we will restrict ourselves to this reading of the term. For a more detailed account of what ontologies are, see [3]. Since ontologies must be not just computer readable but also computer interpretable, an ontology is more than a list of terms stored in a computer parsable format; it comprises also the semantic content associated with these terms – or at least it is supposed to do so.

Before we take a closer look at how this works, we need to make some terminological distinctions. First, we shall use the term ‘type’ in what follows to refer to those entities in reality which terms in ontologies designate. Second,

*This is a preprint of the chapter ‘Modelling Principles and Methodologies – Relations in Anatomical Ontologies’ in Albert Burger, Duncan Davidson and Richard Baldock (Editors): *Anatomy Ontologies for Bioinformatics: Principles and Practice* (to appear in 2007). This work was funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant 1 U 54 HG004028. Information on the National Centers for Biomedical Computing can be found at: <http://nihroadmap.nih.gov/bioinformatics>

it is important that we distinguish between terms (like ‘female pelvis’) and ‘*part_of*’ on the one hand, and the semantic content of such terms on the other. In biomedical ontologies there are two kinds of terms: those denoting types (to be more specific, biomedical types) and those denoting relations. Thus in an anatomy ontology the term ‘female pelvis’ denotes the type *Female Pelvis* and the term ‘*part_of*’ denotes the parthood relation. (Throughout this chapter, we use italics and initial capitals when using type terms to denote types, and quotation marks when we need to talk about the terms themselves. Further we will assume that all terms denote types of the human anatomy, if not explicitly stated otherwise.) Terms are linguistic entities that are created by humans; they satisfy linguistic conventions created by humans; and they can be used to create sentences that express statements about the world. It is terms that are the bearers of semantical content, which means they have denotations – which for present purposes are types and relations.

Note that the types and relations that we are talking about are not the familiar entities that we know from set theory. In [10] we propose a distinction between types (for example the type *Ear*, which is what particular ears share in common) and the sets which are the extensions of such types (the collection of all particular ears). This distinction should be borne in mind to avoid certain sorts of confusion. The types of the biomedical domain (also sometimes called ‘universals’ or ‘kinds’) are the patterns in reality that scientists study and describe in their theories.

Sets and types behave similarly in one important respect: just as sets have *members*, so types have *instances*. However, there are important differences between the two. First: for any arbitrary chosen group of individuals there is a corresponding set, but there need not be a corresponding type. For example, there is the set whose members are exactly: Barbara Bush, Bill Clinton’s left foot, and a given red blood cell of my dog; but there is no corresponding type in reality of which exactly these entities are instances. Types are contrasted with such arbitrary collections by the fact that they can serve as objects of scientific investigation and play a role in scientific generalizations (some of which are then captured in ontologies). A second important distinction turns on the fact that the membership relation is timeless, whereas the instantiation of types is time-dependent. For example, an animal that instantiates the type *Adult Frog* now used to instantiate the type *Tadpole* at some time in the past. Hence types like *Adult Frog* and *Tadpole* gain and lose instances over time, where sets cannot gain or lose members. (If you ‘add’ a member to a set, then the result will be a different set.) Individuals, similarly, can gain and lose parts. However individuals, like this blood cell or that heart are distinguished from the types *Blood Cell* and *Heart* by the following criterion: At any time of its existence an individual necessarily occupies a unique spatial location; a type, in contrast, can be (through its instances) fully present at multiple locations. For example you are necessarily present at exactly one location, whereas the type *Humanity* is currently located at about six billion different locations.

With this background it is easy to formulate the problem that ontologies address. As mentioned above, an ontology is not just a list of syntactical strings

like ‘pelvis’, ‘urinary bladder’, ‘body’ (its type terms); its goal is to comprehend also the semantical content of these terms – that is, the types which they denote – in a machine readable form. This is more problematic than one might think. If humans do not understand the meaning of a term, we can use dictionaries. For example, assume that you don’t speak German and you are wondering what the term ‘Handwurzelknochen’ means. If you look it up you will come to the following conclusion:

1 *The German term ‘Handwurzelknochen’ denotes the type Carpal Bone.*

Note that in (1) the German expression is in quotation marks, whereas the corresponding English term on the right hand side is not. This is because in (1) we are using the English term in order to explain the semantic content of the corresponding German term. This strategy works very well – at least for those who understand the expression ‘carpal bone’. But imagine that a young child were to ask you what ‘Handwurzelknochen’ means. Because the child does not know what a carpal bone is, (1) would not be very helpful. In the best case the child would memorize (1) and would afterwards be able to say ‘carpal bone’ whenever somebody asks what ‘Handwurzelknochen’ means. But obviously she would not know the semantic content of either term. For this reason (1) would not be an appropriate answer in the given case. It would be better to explain the term to the child for example with the help of pictures in an anatomy textbook. But while we can explain the semantical content of terms to people with the help of examples, paraphrases, pictures, and translations into other natural languages, these strategies won’t work for computers. A computer has no better understanding of the term ‘carpal bone’ than of ‘Handwurzelknochen’, so statement (1) will provide the computer with no assistance at all in grasping the semantical content of the latter. Of course we could create a digital dictionary which links ‘Handwurzelknochen’ to the term ‘carpal bone’. Such a dictionary might be useful, because it allows human users to find the appropriate translation; but in this case the computer would be in a similar situation as the child who knows that the two terms have the same meaning, but does not understand either of them. And while we can use pictures to educate the child, this strategy, too, will not work for computers. For them we need an alternative approach – ontologies.

In a first, rough formulation, the main idea of ontologies is the following. The semantical content of expressions that denote types is captured in ontologies through the assertion of relations between the types. This can be easily seen if we conceive an ontology as a graph, whose nodes are labeled with type terms like ‘myelin’ or ‘lipoprotein’ and whose edges are labeled with relation terms like ‘*is_a*’, ‘*part_of*’ and ‘*located_in*’. A graph with many labels is a syntactic entity, but it is important to notice that in an ontology each edge of the graph is equivalent to a statement about the corresponding entities in the biomedical domain. For example, if a node labeled ‘myelin’ and another node labeled ‘lipoprotein’ are connected by an edge labeled ‘*is_a*’, then the ontology expresses the statement ‘*Myelin is_a Lipoprotein*’. Hence the ontology contains claims about the relations that hold between the types that are the denotations of the

type terms of the ontology – and it is here that the ontology gains semantic traction.

It is important to notice that in a well-constructed ontology the type terms will be connected by many links to other terms, thus creating a semantical network each link of which represents a statement about some ontological relation between the types in reality represented by its nodes. The idea is that the semantic content of a term is not determined by any one specific link, but rather by its connection with many other terms that charges it with semantic content. This holistic approach to semantics is not new; it is a central feature of de Saussure's structuralism [6] or of Trier's word field theory [12]. Both de Saussure and Trier identified the meaning of a term with its position in a semantic network of terms. Note however that the holistic thesis in this radical form, although accepted by many contemporary computer science ontologists, is not plausible. For consider the following statements:

2 *Shkart is-a Trkarp.*

3 *Brajhn is-a Trkarp.*

4 *Trkarp part-of Xriprg.*

If the semantical content of a type term in an ontology were completely determined by its relation to other terms, we would have a good understanding what 'trkarp' means by looking at (2-4). However, no semantical content is fixed by (2-4), and even adding further similar statements would not bring about a change in this respect. This does not mean that additional statements would not make a difference. On the contrary: each of them puts additional restrictions on the use of the corresponding terms, and thus helps to narrow down their semantic content. For example, (2-4) allow for the possibility that 'trkarp' denotes *Pelvis*, 'shkart' denotes *Female Pelvis*, 'brajhn' denotes *Male Pelvis*, and 'xriprg' denotes *Body*. However, if we add the additional statement (5), then this possibility is eliminated.

5 *Shkart is-a Brajhn.*

Imagine we were to add hundreds of additional statements to our list by using the terms 'shkart', 'trkarp', 'brajhn', 'xriprg' together with a few dozen other similar fantasy type terms. Each connection in the resultant semantic network would restrict the possible interpretations of 'trkarp' and the other terms and thus provide us with extra semantical content. However, even with hundreds of additional statements it would still not be possible to determine which type is denoted by 'trkarp', for there will still be many possible interpretations left. The meaning of the terms will thus not be completely determined, and so the holistic thesis in its radical version is false.

It is important to realize the falsity of radical holism, because this has an important consequence: since an ontology can't completely fix the semantics of a type term, the semantic content of a term is in this sense not an all-or nothing matter but a matter of degree: the more a term is connected to other terms,

the more its semantical content is determined. A sparse ontology that consists of only loosely connected terms provides these terms with very little in the way of semantic content. In particular, type terms in an ontology that are not distinguished by their connections within the network of terms are semantically indistinguishable with respect to that ontology. For example, assume that we have an ontology that consists of (2-4) and (6).

6 *Shkart and Brajhn are disjoint.*

Statement (6) guarantees that ‘shkart’ and ‘brajhn’ do not denote the same type. However, even with (6) the terms ‘shkart’ and ‘brajhn’ would still not be distinguishable: all we know about them is that they denote subtypes of *Trkarp* and that their denotations are disjoint. This limits the value of the given ontology for applications. For example, assume that two scientists use the ontology to annotate their data and that one of them believes that ‘shkart’ denotes *Male Pelvis* while the other believes that it denotes *Female Pelvis*. Since the ontology does not contain any information about the difference between male and female pelvises, an automatic reasoner would never be able to detect that the scientists are using the term ‘shkart’ in a crucially different way. This example shows why sparse ontologies are inferior to rich ontologies; the latter convey a greater amount of semantic content.

3 The semantic content of relation terms

Let us recap the results so far. An ontology is more than a list of type terms, it is designed to encapsulate also the ‘meaning’ of these terms in a computer parsable form. Biomedical ontologies consist of statements that involve type terms and a relation term; since ontologies are often visualized as graphs, it is helpful to think of their type terms as labels attached to nodes and of the relation terms as labels attached to edges. Since there is no way to tell a computer directly which type is determined by a given type term, ontologies seek to do this indirectly. The basic idea is that the semantic content of a type term is captured by its position in the network of type terms of which it is a constituent. Since each statement expresses a relation between the denotations of its type terms, each statement limits the possible interpretations of its type terms. For example, (7) expresses the thesis that the denotations of ‘trkarp’ and ‘xriprg’ are related by the parthood relation, thus limiting the possible interpretations of ‘trkarp’ and ‘xriprg’.

7 *Trkarp part_of Xriprg.*

The approach will not be sufficient to single out some specific type as denotation of each given type term, but if the term is connected to a multitude of other terms, then the possibilities will be correspondingly restricted. Note that, according to this approach, the semantical content is determined by restricting the possible interpretations of the type terms via the relations between their respective denotations – in our example the parthood relation between ‘trkarp’

and ‘xriprg’. However, while humans know that the relation term ‘*part_of*’ is supposed to express the parthood relation, computers do not. For a computer ‘*part_of*’ is just a string like any other, and for the computer (7) is itself a string which is not intrinsically different from a string such as (8):

8 *Trkarp caxc Xriprg.*

How, then, do we bridge the gap between relation terms and the relations themselves? When we describe the links between the types we use relation terms like ‘*is_a*’ and ‘*part_of*’; but how do we capture the denotations of the latter in a machine interpretable way? This question is important for two reasons. First, it is relations which form the principal vehicle for interoperability of ontologies. Thus if the same relations can be used in all members of a given set of ontologies, then to this degree these ontologies form an interoperable family – an idea which forms one central pillar of the OBO Foundry initiative (<http://www.obofoundry.org>). The use of common relations when creating a system of ontologies is equivalent to the use of a common gauge when creating an international railway system.

Second, the relation terms used in ontologies typically denote rather abstract relations. If we use expressions like ‘*part_of*’, ‘*located_in*’, and ‘*develops_from*’ as unanalyzed primitives, these expressions are semantically underspecified. As shown in [2] and [11], the result is that they are used in an ambiguous way. For example, in the FMA we find:

9 *Female Pelvis part_of Body.*

10 *Urinary Bladder part_of Female Pelvis.*

11 *Urinary Bladder part_of Body.*

Statement (9) is used to assert that every female pelvis is part of a human body, but it does not imply that every body has a female pelvis as part. In contrast, (10) is used to assert that every female pelvis has a urinary bladder as a part, but not that every urinary bladder is part of a female pelvis. The parthood relation between the types denoted in (11) is the strongest of the three: Every urinary bladder is a part of a body and every body has a urinary bladder as part.

Another example is the use of ‘contains’ in GALEN, where we find:

12 *Pelvic Cavity contains Ovarian Artery.*

13 *Male Pelvic Cavity contains Urinary Bladder.*

14 *Tooth Socket contains Tooth.*

Here the different statements express very different states of affairs, because the relation term ‘contains’ is used ambiguously. For every ovarian artery there is a pelvic cavity such that the pelvic cavity **contains** the ovarian artery. However, not every pelvic cavity **contains** an ovarian artery. This is expressed by

(12). In contrast (13) states that every male pelvic cavity **contains** a urinary bladder, but it does not say that every urinary bladder is contained in a male pelvic cavity. In (12) and (13) ‘contains’ denotes distinct relations holding, respectively, between a type of immaterial entity (a cavity) and types of material objects (arteries, urinary bladders). In both cases the material objects are completely located in the cavities. In contrast, ‘contains’ in (14) relates two types of material objects (tooth sockets and teeth). Further the teeth are only partially contained in the tooth sockets. Hence ‘contains’ in (14) expresses a relation that is very different from the relations expressed by the same term in (12) and (13).

The fact that ‘*part_of*’ and ‘*contains*’ in statements (9-14) are used ambiguously would be less problematic if the statements were to appear in a text that is intended to be read by humans with some knowledge of anatomy. A human can use background knowledge to disambiguate the statements in appropriate ways. However, a computer is not able to handle ambiguity in the way a human can, so that it is crucial for an ontology that relation terms are used in a clear-cut way; otherwise automatic reasoning is bound to lead to false conclusions (see [2] for examples). In addition, since the relations are used in ontologies to determine the semantical content of the terms in the ontology, a lack of clarity with respect to the relations will contaminate the whole ontology. For this reason, too, therefore it is essential for the use of an ontology that the semantics of the relation terms be made explicit in a non-ambiguous way.

The first step in solving this problem is to distinguish between relations that hold between types and those that hold between the instances of those types. Ontologies are about types: Statement (9) asserts that a parthood relation holds between the type *Female Pelvis* and the type *Body*, statement (12) that a containment relation holds between the type *Pelvic Cavity* and the type *Ovarian Artery*, etc.

Since the type terms in an ontology denote types and the relation terms like ‘*part of*’, ‘*is a*’, ‘*develops from*’ denote relations between types, instances might seem to be not important for an ontologist. However, an anatomist is not able to study the types directly. We have epistemic access to types only via their instances. Hence the only way to evaluate a statement concerning types – for example ‘*Pelvic Cavity contains Ovarian Artery*’ – is to look at instances of *Ovarian Artery* and their locations; there is no way to look at the type *Ovarian Artery* directly. Similarly, the only way to evaluate a statement like ‘*Appendix part_of Body*’ is to look at instances of the type *Appendix* and to check whether they are part of some instances of the type *Body*. Note that the parthood relation between the instances differs from the various parthood relations on the type-level that we have considered above. One major difference is that the parthood relation between anatomical structures (i.e. between the different types of anatomical structures) holds in a time-dependent way. For example, it might be the case that Bill’s appendix is part of his body at 6 am, but that it is not part of his body at 8pm on the same day. The relation expressed in ‘*Appendix part_of Body*’ is however a timeless relation between types. Arguably, another difference is that the fact that Bill’s appendix is part

of his body at a given time entails that the location of his appendix and the location of his body overlap at this time (where it is not clear what it would mean for the type *Appendix* to have a location that overlaps with the location of the type *Body*). To capture the differences we will henceforth distinguish relations between types, for which we use *italic font*, from relations of other kinds, picked out by using **bold**.

We begin with studying the **part_of at** t relation between instances (where t stands in for times). Only by studying this and similar relations on the instance level can we gain insight into the parthood relations on the type-level ([2], [8]). Although the latter are at first importance for ontologies, they are actually secondary to the former in an epistemic sense. We can use the tight connections between **part_of at** a given time and the parthood relations on the type-level to disambiguate the use of ‘*part_of*’ in the problematic cases mentioned above. In cases (9-11) the relation term ‘*part_of*’ can be read as denoting three different relations; hence we have to distinguish at least three different parthood relations that hold between types. In the following we will use the term ‘*part_of*’ only to denote one of these relations; for the others ones we will use the terms ‘*is_part*’ and ‘*integral_part_of*’.

Let C and C_1 be types of anatomical entities, let x, y, z be anatomical entities (instances), and t a time. Further, let ‘ Cyt ’ be the abbreviation for ‘ y is an instance of C at time t ’ and, ‘ C_1zt ’ the abbreviation for ‘ z is an instance of C_1 at t ’. We can now define:¹

d 1 $C \text{ part_of } C_1 =_{def}$ for all y, t , if Cyt then there is some z such that C_1zt and $y \text{ part_of } z \text{ at } t$.

d 2 $C \text{ is_part } C_1 =_{def}$ for all z, t , if C_1zt then there is some y such that Cyt and $y \text{ part_of } z \text{ at } t$.

d 3 $C \text{ integral_part_of } C_1 =_{def} C \text{ part_of } C_1$ and $C \text{ is_part } C_1$.

These definitions provide an example of how we can define relations between types in terms of the relations between the corresponding instances. One major advantage of these definitions is that they provide us with a better understanding of the type-level statements that form an ontology. With the help of the definitions (d 1 - d 3) it is easy to see that (15) is true, but (16) false, in virtue of the fact that there are (human) bodies that have no female pelvis (because they have a male pelvis).

15 *Female Pelvis part_of Body.*

16 *Female Pelvis is_part Body.*

¹The terms ‘*part_of*’ and ‘*integral_part_of*’ are defined as in [8], *is_part* is the inverse of *has_part* as defined in [8], which means that $C \text{ is_part } C_1$ is logically equivalent to $C_1 \text{ has_part } C$. The relations *part_of*, *is_part*, and *integral_part_of* are equivalent to P_1 , P_2 , and P_{12} as defined in [2] and in [1] in this book.

In addition, the definitions (d 1 - d 3) allow us to check the logical properties of the type level relations and their logical connection. Without the definitions it might not be obvious whether ‘A *part_of* B’ implies ‘B *is_part* A’ and vice versa, or in other words whether *part_of* is the inverse of *is_part*. With the help of the (d 1 - d 3) it is easy to see that this is not the case.

Let us consider another example. Does (15) entail (17)?

17 *Body is_part Female Pelvis.*

According to the definition (d 2) the statement (17) means: For any instance of *Female Pelvis* at any time, there is some instance of *Body* such that that instance of *Body* is part of that instance of *Female Pelvis* at that time. Since human bodies are never parts of pelvises, this is obviously false – hence we have shown that *part_of* is not the inverse of *is_part* ([2], [8]).

Let’s consider two other examples. Since men have urinary bladders, some urinary bladders are not part of a female pelvis. Hence (18) is false. In contrast, (19) is true, because female pelvises have urinary bladders as parts:

18 *Urinary Bladder part_of Female Pelvis.*

19 *Urinary Bladder is_part Female Pelvis.*

The definitions (d 1 - d 3) provide us with a clear understanding of the relations which allows us to use the corresponding assertions to draw logical inferences. To give a very primitive example, from (d 1 - d 3) it follows immediately that (20) entails (21) and (22). Such logical connections facilitate automatic reasoning (see [1] in this book).

20 *Urinary Bladder integral_part_of Body.*

21 *Urinary Bladder part_of Body.*

22 *Urinary Bladder is_part Body.*

In the beginning of this section we addressed two problems: (a) Humans use relation terms like ‘*part_of*’ ambiguously, which undermines the quality of ontologies and leads automatic reasoners astray. And (b) the relations are used in ontologies to determine the semantic content of the type terms, hence we need to capture the denotation of relation terms like ‘*part_of*’ in a machine interpretable form. These problems were addressed by defining relations on the type-level (in our example *integral_part_of*, *part_of*, and *is_part*) with the help of a relation between individuals (**part_of at**). The definitions (d 1 - d 3) allow us to resolve the ambiguities in existing uses of the term ‘*part_of*’ and it is easy to translate the definitions above into a formal language, hence the approach that was embraced in the last section was a step in the right direction.

However, it did not solve the problems completely. The definitions (d 1 - d 3) involve the parthood relation between individuals. Hence the denotation of the terms ‘*part_of*’, ‘*is_part*’, and ‘*integral_part_of*’ depends on the denotation

of the term ‘**part_of at**’ that we used in these definitions. Thus in order to get a clear understanding of these terms we need to determine the denotation of ‘**part_of at**’. This can be done via an axiomatization of this relation, i.e. by providing a set of axioms which amount to a so-called ‘contextual definition’ of ‘**part_of at**’. A contextual definition is not really a definition in the strict sense, but the axioms serve to capture our intuitions about the logical properties of the relation that is axiomatized and thus they restrict the possible interpretations of the term ‘**part_of at**’.

It would have been possible to axiomatize the various parthood relations on the type-level directly instead of defining them with the help of the parthood relation on the instance-level. There are however two reasons why it is better not to do this, but to use the **part_of at** relation as we have done. One reason is that we could use the single parthood relation **part_of at** on the level of individuals to define the three relations on the type-level. Thus we needed only one primitive notion instead of three. Further, since we have access to types only via their instances, our intuitions about the logical properties of the relations on the instance-level are much more developed. Moreover, much of our digital data about anatomical and other entities in the biomedical domain comes in the form of the instance data contained, for example, in clinical records.

Actually, since ‘part’ is not a technical term but an expression we use in daily life (we talk about engine parts, or the parts of former Yugoslavia, or about cellulose as part of wood) one might suspect that we have very strong intuitions about the parthood relation and thus that it would be easy to develop a theory of wholes and their parts. Indeed it is true that people have strong opinions on mereological questions; unfortunately the intuitions governing our daily talk about wholes and their parts are quite heterogeneous (if not plainly inconsistent). For this reason mereology is a controversial field in philosophy. Hence it is important to give an explicit account of **part_of at**, otherwise type-level terms like ‘*integral_part_of*’, ‘*part_of*’, and ‘*is_part*’ will themselves be used ambiguously. This is not the place to present a full axiomatization (see [7]), but some examples of axioms that many people would embrace are:

Ax. 1 *At any time t , every x **part_of** x **at** t .*

Ax. 2 *For any x, y, t : if x is **part_of** y **at** t and y **part_of** x **at** t , then $x = y$.*

Ax. 3 *For any x, y, z, t : if x is **part_of** y **at** t and y is **part_of** z **at** t , then x is **part_of** z **at** t .*

These axioms express time-relativized versions of the reflexivity, antisymmetry, and transitivity of parthood, respectively.

The approach that we have considered in this section allows us to restrict the semantical content of relation terms. This is achieved in two steps. We define the type-level relation with the help of a relation between individuals and we then give an axiomatization of the latter relation. This approach has been presented by means of by appealing to just a few examples and is still rather sketchy. In [1] in this book it will be covered systematically and in greater depth.

4 Canonicity

So far we did not discuss one important objection to the above approach.² Let's assume that we encounter statement (23) in a textbook on human anatomy.

23 *Appendix is-part Body.*

Statement (23) is true: the human body has an appendix. However, according to the definition (d 2) statement (23) is equivalent to (24):

24 *For every x and time t , if x is a Body at t , then there is an instance of Appendix y at t such that y is **part_of** x at t .*

Statement (24) is plainly false: there are plenty of people who live happily without an appendix. Each of them provides a counterexample to the claim in (24) that every body has an appendix. Since (23) is true, but (24) is false, the statements (23) and (24) cannot be equivalent. Does that mean that our analysis of statements like (23) is incorrect? Is definition (d 2) inappropriate?

In order to understand the root of the problem we need to distinguish between canonical anatomy and instantiated anatomy ([5], [9]). Instantiated anatomy concerns the anatomical entities represented for example in data about actual cases generated in clinical practice. Canonical anatomy is the result of generalizations deduced from qualitative observations that are implicitly sanctioned by their accepted usage by anatomists. While instantiated anatomy and canonical anatomy are both founded in empirical observations, only instantiated anatomy contains empirical statements about human bodies and their anatomical parts. In contrast, the relation between canonical anatomy and human bodies is in some respects similar to the relation between a technical drawing and the artifacts that are built with the help of the drawing. As anybody who has assembled a piece of Swedish furniture knows, many existing artifacts do not exactly match their technical drawings. That does not make the technical drawing 'false'; a technical drawing is not an empirical description of the composition of the existing artifacts; rather it tells us how the artifacts should be composed. Analogously, a canonical anatomy gives an account of the 'prototypical' composition of the male or female human body. For example, (23) does not assert that all human bodies have an appendix, but rather that a human body is supposed to have an appendix. Thus (23) cannot be refuted by the fact that some people lack an appendix. This example shows that a canonical anatomy consists of statements that describe how the anatomical entities of a given organism are supposed to be composed (for example in light of the structure of the underlying genes); and it is this that distinguishes a canonical anatomy from an instantiated anatomy.

The distinction between instantiated and canonical anatomy is important since it allows us to analyze the source of the mismatch between (23) and (24). Statement (24) would be an appropriate analysis of (23) if (23) would be an

²We thank Cristian Cocos, Alan Rector, and Cornelius Rosse for their critical remarks and suggestions.

assertion about instantiated anatomy – and in this case (23) would be false, since it is an empirical fact that not all human bodies have an appendix. However, we have assumed above that (23) is a statement within a textbook on canonical anatomy. One way to make the force of statements of this kind explicit is to use an adverb as in (25):

25 *Canonically, Appendix is_part Body.*

Syntactically, the expression ‘canonically’ in (25) works like ‘necessarily’, ‘possibly’, ‘it is permissible that’ and other expressions that are – from a logical perspective – logical operators. However, while the semantics of the latter is well understood, the semantics of ‘canonically’ is not. Thus in this form (25) is a logical black box and for this reason useless for logical reasoning. This is why we will present a logical analysis of statements like (25) in the remainder of this section.

In [8] we have (implicitly) embraced the assumption that in the context of canonical anatomy the domain of discourse is restricted to canonical entities. In this case (25) would have the same meaning as (23), except for an implicit understanding that we consider only canonical entities – which can be made explicit by restricting the range of the variables in (24). Hence – according to this approach – (25) is equivalent to (26):

26 *For every x and time t , if x is a Body at t , then there is an instance of Appendix y at t such that y is **part_of** x at t ; where the variables x and y range exclusively over canonical entities.*

The term ‘canonical entity’ can be defined as follows:

d 4 *An anatomical entity x is canonical with respect to a given anatomy A if and only if x is structured in the way it is supposed to be structured according to anatomy A .*

Since a human body without an appendix is not canonical, it follows that such bodies fall outside the domain of quantification, and thus the problematic cases are excluded.

Unfortunately, this way of understanding ‘canonically’ leads to new difficulties. For example, (27) would be equivalent to (28).

27 *Canonically, Appendix part_of Body.*

28 *For every x and time t , if x is an instance of Appendix at t , then there is an instance of Body y at t such that x **part_of** y at t , where the variables x and y range exclusively over canonical entities.*

Statement (28) expresses that every canonical appendix is part of a canonical body – which is obviously wrong: there are people who have a perfectly normal appendix, but are lacking teeth. Therefore the idea of restricting the domain of quantification to canonical entities does not work; we need to find an alternative way to analyze (25).

In order to come up with the needed analysis, we have to remember that canonical anatomy gives an account of how a male or female human body is supposed to be composed. Thus a statement that is part of a canonical anatomy expresses a requirement that a human body has to meet in order to conform to the given canonical anatomy. We can express this in the following way: Let ‘ I_Axt ’ be the abbreviation for ‘ x is a human body that is in conformity with anatomy A at t ’. (As mentioned above we assume that we deal with human anatomy; otherwise one has to modify the definition of ‘ I_Axt ’ in the obvious way.)

d 5 *Canonically, C is_part $C_1 =_{def}$ for all x, t , necessarily, if I_Axt , then FOR ALL z , IF C_1zt THERE IS SOME y SUCH THAT Cyt AND y **part_of** z **at** t ; where y and z are anatomical entities that are **part of** x **at** t .*

Definition (d 5) can be paraphrased as follows: if a statement of the form ‘ C is_part C_1 ’ is part of a canonical human anatomy, then the following holds for any human body x at any given time: necessarily, if x is in conformity with the given anatomy (at this time), then, for any anatomical part of x that is an instance of C_1 (at this time), there is an anatomical part of x that is an instance of C (at this time) and the instance of C is part of the instance of C_1 (at this time).

Let’s consider an example. Definition (d 5) entails that (29) is equivalent to (30):

29 *Canonically, Carpal Bone is_part Hand.*

30 *Necessarily, if x is a human body that is in conformity with A at time t , then for all y , if y is an instance of Hand at t , there is (at least) one entity z that is an instance of Carpal Bone at t and is **part of** y at t ; where y and z are anatomical entities that are **part of** x **at** t .*

Analogously, we can define *part_of* for canonical anatomies:

d 6 *Canonically, C part_of $C_1 =_{def}$ for all x, t , necessarily, if I_Axt , then FOR ALL y, t , IF Cyt , THERE IS SOME z SUCH THAT C_1zt AND y **part_of** z **at** t ; where y and z are anatomical entities that are **part of** x **at** t .*

Definition (d 6) expresses the following: if a statement of the form ‘ C part_of C_1 ’ is part of a canonical human anatomy, then the following holds for any human body x at any given time: necessarily; if x is in conformity with A (at this time), then, for any anatomical part of x that is an instance of C (at this time), there is an anatomical part of x that is an instance of C_1 (at this time) and the instance of C is part of the instance of C_1 (at this time).

The definitions (d 6) and (d 5) are closely linked to (d 1) and (d 2), respectively: the parts of the definitions that are emphasized by using small caps are the right hand sides of the definitions (d 1) and (d 2). We chose this way of presenting the definitions because it shows that the original account of the last section is preserved, it is just that it is now embedded in a context that does justice to the fact that statements like (23) are part of a canonical anatomy.

Let's consider another relation, where time plays a more important role than in the examples above. The human body is supposed have deciduous teeth and the human body is supposed to have androgenic hair – but obviously not at the same time. We can express this fact with the help of a relation *excludes* in (31), where *excludes* is defined in definition (d 7).

31 *Canonically, Deciduous Tooth excludes Androgenic Hair.*

d 7 *Canonically, C excludes $C_1 =_{def}$ for all x, t , necessarily, if $I_A x t$, then THERE ARE NO $y, z t$, SUCH THAT $C y t$ AND $C_1 z t$; where y and z are anatomical entities that are **part of x at t** .*

The relation *excludes* serves here as a simple example that illustrates how time can play an important role for the definitions of type-level relations; this holds in particular for relations that concern the development of anatomical entities.

Since we have focused on parthood relations in this section so far, let's consider an example that involves the *contains* relation between types, e.g. (32). Further, let's assume that we have an account of the corresponding **contains** relation on the instance-level (see [1]). We can now define *contains* with the help of **contains** as in definition (d 8).

32 *Canonically, Male Pelvic Cavity contains Urinary Bladder.*

d 8 *Canonically, C contains $C_1 =_{def}$ for all x, t , necessarily, if $I_A x t$, then FOR ALL y , IF $C y t$ THEN THERE IS SOME z SUCH THAT $C_1 z t$ AND y **contains z at t** ; where y and z are anatomical entities that are **part of x at t** .*

Hence (32) is equivalent to (33), which is itself a complicated way of expressing (34):

33 *Necessarily, if $I_A x t$, then for all y , if y is an instance of Pelvic Cavity at t then there is some z such that z is an instance of Urinary Bladder at t and y **contains z at t** ; where y and z are anatomical entities that are **part of x at t** .*

34 *Necessarily, if x is a human body that is in conformity with A at t , and x has a pelvic cavity, then there is a urinary bladder that is contained in the pelvic cavity.*

We will now generalize our approach and define 'canonically'. In this section we have analyzed statements of the form (35), where *rel* stands in for '*is_part*', '*part_of*', '*excludes*', and '*contains*'.

35 *Canonically, C rel C_1 .*

As we have mentioned above, the definitions (d 5) and (d 6) (where *rel* is *is_part* and *part_of*, respectively) are closely linked to the definitions (d 1) and (d 2), which define statements of the form '*C is_part C₁*' and '*C part_of C₁*'. Analogously, the definition (d 8) is closely linked to (d 9). (Again, the relevant parts of the definition (d 8) are in small caps.)

d 9 C contains $C_1 =_{def}$ if Cyt , then there is some z such that C_1zt and y contains z at t .

It seems that for any definition that defines statements of the form (35), there is a corresponding definition of the statements that does not begin with ‘canonically’. We will use this connection in order to define ‘canonically’:

d 10 Let rel be any binary type-level relationship, and C, C_1 any types, and assume we have a definition of the following form:

$$C \text{ rel } C_1 =_{def} \phi(y, z)$$

where $\phi(y, z)$ represents a formula that involves only relationships between individuals and that ensures that y and z are anatomical entities that are instances of C and C_1 , respectively. In this case we can define:

Canonically, $C \text{ rel } C_1 =_{def}$ for all x, t , necessarily,

if I_Axt , then $(\phi(y, z)$ and y **part of** x **at** t and z **part of** x **at** t)

Definition schema (d 10) provides us with a systematic link between the relations within a canonical anatomy and the use of the corresponding relations within an instantiated anatomy. The definition schema (d 10) works not only for relations such as those that we have considered in this section; it can be applied to many type-level relations and in particular to the type-level spatial relations that will be considered in [1].³

5 Conclusions

One purpose of an ontology is to encapsulate the meanings of its terms in a computer parsable form. We analyzed how anatomical anatomies fulfill this purpose. An anatomical ontology consists of statements composed of two kind of terms denoting types and relations, respectively. Typically such statements involve two type terms, so that they are of the form ‘ $A \text{ rel } B$ ’.

We showed that there is no way to tell a computer directly, for any given type term, which type is denoted by that term. Thus ontologies must find ways to convey such information indirectly: broadly, it is the totality of the relations between the denotations of the type terms that determines the semantical content of the type terms taken individually. This works as follows. Each statement

³Note that the connection between the relations is not always as straight forward as in the examples considered above. For example, definition (d 11) does not capture the semantic content of ‘excludes’.

d 11 C excludes $C_1 =_{def}$ there are no y, z, t , such that Cyt and C_1zt .

A more appropriate definition of ‘excludes’ is:

d 12 C excludes $C_1 =_{def}$ there are no u, y, z, t , such that y **part of** u **at** t , z **part of** u **at** t , Cyt , and C_1zt .

‘ $A \text{ rel } B$ ’ asserts that the denotation of ‘ A ’ and the denotation of ‘ B ’ are linked by the relation *rel*. Thus any interpretation of ‘ A ’ and ‘ B ’ according to which their denotations do not meet this requirement is ruled out. The possible interpretations of the terms ‘ A ’ and ‘ B ’ are in this sense limited by the statement ‘ $A \text{ rel } B$ ’. While this approach is not sufficient to single out any specific type as denotation of a given term, if the term is connected to a multitude of other terms, then the possibilities will be correspondingly restricted. Fortunately, in the domain of anatomy we are already in possession of high-quality representations of such multiple relations.

Since the semantical content of type terms is determined by the relations that are expressed by ‘*is_a*’, ‘*part_of*’, ‘*contains*’ and other relation terms, it is crucial to make explicit which relations these terms denote. This analysis is important not only because of our aim to capture the semantical content of the terms of an ontology in a machine-readable form, but also because people tend to use relation terms ambiguously, in a way which reduces the quality of ontologies. We showed that many relations between types can be defined with the help of relations that hold between instances of these types, and an approach based on this recognition has the advantage that we typically have a better understanding of the relations between instances than of the relations between the corresponding types. Further, the approach has the virtue of economy, since it is often possible to define different relations on the type-level with the help of one relation on the instance-level.

On the given approach the meaning of a statement ‘ $A \text{ rel } B$ ’ in an ontology is an empirical assertion about the instances of types A and B . Thus ‘*Embryo develops_from Zygote*’ is true if and only if: for any instance of *Embryo* x there is an instance of *Zygote* y such that x **developed_from** y . Here **developed_from** is an instance level relation that holds between individuals. ‘*Embryo develops_from Zygote*’ is thus an empirical assertion that can be falsified (by discovering that at least one embryo did not develop from a zygote).

In the case of canonical anatomical ontologies such as the FMA, in contrast, the situation is more complicated, since canonical anatomical ontologies do not consist of empirical assertions in this sense, but rather of statements that express how the corresponding entities are supposed to relate to each other (in virtue of the workings of the underlying structural genes). For this reason we analyzed statements of the form ‘*canonically, A rel B*’ in such a way as to show how the semantic content of such statements is systematically linked to statements without the prefix ‘*canonically*’. Very roughly, a statement like ‘*canonically, A rel B*’ expresses that, necessarily, any human body x that is in conformity with the given anatomy meets the requirement ‘ $A \text{ rel } B$ ’, where ‘ $A \text{ rel } B$ ’ can spelled out as in the non-canonical case and the domain of discourse is restricted to the anatomical entities that are part of x . The fundamental picture then remains the same: the semantic content of the type terms is provided by the network of relations between them. A profound understanding of these relations is thus a prerequisite for a non-ambiguous use of type terms of the sort which can support automatic reasoning. The next chapter will present a deeper and more systematic analysis of those specific sorts of spatial relations that are relevant

for anatomical ontologies.

References

- [1] T. Bittner, M. Donnelly, L. Goldberg and F. Neuhaus. Modelling Principles and Methodologies – Spatial Representation and Reasoning Albert Burger, Duncan Davidson and Richard Baldock (eds.): *Anatomy Ontologies for Bioinformatics: Principles and Practice* (in print).
- [2] M. Donnelly, T. Bittner, and C. Rosse. A formal theory for spatial representation and reasoning in bio-medical ontologies. *Artificial Intelligence in Medicine*, 36(1):1–27, 2006.
- [3] N Guarino. Formal Ontology and Information Systems. In N Guarino (ed.), *Proceedings of the 1st International Conference on Formal Ontologies in Information Systems, FOIS'98*, IOS Press, Trento, 3-15, 1998.
- [4] I. Johansson. On the Transitivity of the Parthood Relations. In H. Hochberg and K. Mulligan (eds.), *Relations and Predicates*, 161-181, 2004.
- [5] C Rosse, JL Mejino , BR Modayur , R Jakobovits, KP Hinshaw, JF Brinkley. Motivation and organizational principles for anatomical knowledge representation: the digital anatomist symbolic knowledge base. *Journal of the American Medical Informatics Association*, 5(1):17-40, 1998.
- [6] F de Saussure. *Grundfragen der allgemeinen Sprachwissenschaft*. Walter De Gruyter, Berlin/Leipzig, 1967.
- [7] P. Simons. *Parts, A Study in Ontology*. Clarendon Press, Oxford, 1987.
- [8] B. Smith, W. Ceusters, B. Klagges, J. Köhler, A. Kumar, J. Lomax, C. Mungall, F. Neuhaus, A. Rector, and C. Rosse. Relations in biomedical ontologies. *Gnome Biology*, 6(5):46, 2005.
- [9] B Smith, A Kumar, W Ceusters, C Rosse. On carcinomas and other pathological entities. *Comparative and Functional Genomics*, vol 6, 7-8, 379-387, 2005.
- [10] B Smith, W Kusnierczyk, D Schober, W Ceusters. Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain. Proceedings of KR-MED 2006, in press.
- [11] B. Smith and C. Rosse. The Role of Foundational Relations in the Alignment of Biomedical Ontologies. In M. Fieschi, et al. (eds.), *Medinfo 2004*, IOS Press, Amsterdam, 444-448, 2004.
- [12] J Trier. *Aufsätze und Vorträge zur Wortfeldtheorie*. Berlin: DeGruyter, 1973.