

Anatomical Information Science

Barry Smith^{1,2}, Jose L.V. Mejino Jr.³,
Stefan Schulz⁴, Anand Kumar², and Cornelius Rosse³

¹Department of Philosophy, University at Buffalo, Buffalo, NY 14260, USA

²Institute for Formal Ontology and Medical Information Science, Saarland University,
Saarbrücken, Germany

³University of Washington, Seattle, USA

⁴Freiburg University Hospital, Freiburg, Germany

phismith@buffalo.edu
mejino@u.washington.edu
akumar@ifomis.uni-saarland.de
stschulz@uni-freiburg.de
rosse@u.washington.edu

Abstract. The Foundational Model of Anatomy (FMA) is a map of the human body. Like maps of other sorts – including the map-like representations we find in familiar anatomical atlases – it is a representation of a certain portion of spatial reality as it exists at a certain (idealized) instant of time. But unlike other maps, the FMA comes in the form of a sophisticated ontology of its object-domain, comprising some 1.5 million statements of anatomical relations among some 70,000 anatomical kinds. It is further distinguished from other maps in that it represents not some specific portion of spatial reality (say: Leeds in 1996), but rather the generalized or idealized spatial reality associated with a generalized or idealized human being at some generalized or idealized instant of time. It will be our concern in what follows to outline the approach to ontology that is represented by the FMA and to argue that it can serve as the basis for a new type of anatomical information science. We also draw some implications for our understanding of spatial reasoning and spatial ontologies in general.

1 The Foundational Model of Anatomy

The Foundational Model of Anatomy (FMA) is a computer-based representation of the entities and relations which together form the phenotypic structure of the human organism [1,2]. It provides a qualitative spatial reference system for the human body that is designed to be understandable to human beings and also to be navigable by computers. It is intended as a general-purpose resource, which can be used by any biomedical application that requires anatomical information, from radiology (in supporting automatic image analysis) to pharmacokinetics (in representing the pathways of drugs as they are absorbed by, distributed through, metabolized in and excreted from the body).

The FMA began its life as a classification of anatomical entities called the University of Washington Digital Anatomist Vocabulary. In recent years it has grown

from a list of terms linked by *is_a* and *part_of* relations to a sophisticated spatial-structural ontology of the human organism at all biologically salient levels of granularity, comprehending some 1.5 million statements of ontological relations among some 70,000 anatomical universals. The acronym 'FMA' is currently used in the biomedical informatics community both for this ontology and also for its representation in computerized form within the Protégé 2000 frame-based ontology editing environment [2,3].

We shall argue in what follows that the FMA provides a starting-point for a new type of anatomical information science, representing a new application domain with potentially valuable implications also for other branches of Spatial Information Theory.

2 Types of Relations

The FMA relates exclusively to continuant entities (i.e. to entities, such as molecules, cells, lungs, which endure through time while undergoing changes of various sorts) [4]. The Structural Informatics Group at the University of Washington, which developed and maintains the FMA, has itself initiated work on two complementary ventures, called PRO and PathRO – for 'Physiology' and 'Pathology Reference Ontology', respectively [5] – which deal with those occurrent *processes* in which the anatomical entities at different levels of granularity participate. Here, however, we shall concern ourselves exclusively with continuant entities, which exist at the level of particulars or tokens (having determinate spatial locations at each specific point in time) as instantiations of certain corresponding universals or types (kinds, classes).

We can distinguish a number of distinct types of relations between continuant universals which are employed in the construction of an ontology like the FMA [3,6]:

- 1) *is_a* relations, linking one universal to another (more general) universal in a subsumption hierarchy; examples: *liver is_a organ*, *lacrimal lake is_a anatomical cavity*
- 2) static physical relations between continuant universals; examples: *lobe of liver part_of liver*, *nuclear membrane adjacent_to cytoplasm*
- 3) relations between universals instantiated at different stages in the development of an organism; examples: *zygote derives_from ovum*, *adult transformation_of child*.

Cross-cutting all of these are distinctions between:

- a) instance-level relations (such as the parthood relation between your left thumb and your left hand), which obtain between instances of anatomical universals within the canonical organization of the human body;
- b) relations involving also non-canonical anatomical instances including instances of pathological anatomical universals such as *wounded knee* or *amputation stump*;
- c) relations involving entities (*implants*, *food*, etc.) imported into the human body;

- d) relations involving entities (*biopsied samples, excreta, etc.*) exported from the human body.

The FMA itself focuses on relations of types 1) and 2) under heading a). In what follows we expand our scope to include also relations of other types, drawing on recent work, summarized in [7, 8], involving not only the FMA's developers but also representatives of other influential research groups in biomedical ontology.

3 Canonical Anatomy

The term 'anatomy' is used to refer both to anatomical *science* and to that anatomical *structure* which this science describes, a certain ordered aggregate of material objects and physical spaces filled with substances (such as blood) which together constitute a biological organism [2]. In the case of the FMA the structure in question is what is called the 'canonical' structure of the adult human body, whereby the idea of canonicity (first proposed for the FMA in [9]) has no analogue in geospatial science. For where geospatial maps deal in every case with specific instances (with specific portions of the surface of the earth), the FMA deals not with the *instances*, the individual human beings, whose bodily organization has been investigated over the centuries with the aid of surgical dissection, radiological imaging and other techniques, but rather with a certain ('canonical') idealization thereof (actually with two idealizations, corresponding to the male and female adult human beings, respectively). The FMA, that is to say, is a collection of generalizations deduced from the qualitative observations of the *normal* human body, generalizations which have been refined and sanctioned by successive generations of anatomists and presented in textbooks and atlases of structural anatomy. One needs to take such an idealization as target in a venture like the FMA since the effort to do justice to anatomical structure in all its variants and instantiations would, in the absence of such an idealized reference frame, give rise to an endeavor of unmanageable complexity.

4 Boundaries

A further apparent distinction between the geospatial and anatomical domains from an ontological perspective turns on the fact that, where anatomy embraces within its purview primarily three-dimensional entities such as cells, organs and whole organisms, geospatial ontologies are focused on the (broadly) two-dimensional entities that form the surface of the earth. Applications of geospatial reasoning have thus far been correspondingly concerned with the movement of objects across this surface and with associated questions of land-use, soil-type, forest-coverage, and so forth. Closer inspection reveals, however, that both anatomical and geographic information sciences must deal with entities in all spatial dimensions. Thus the FMA deals not only with material objects but also with both fiat and bona fide boundaries of two, one and zero dimensions [10,11], and a geospatial ontology like that which underlies the Spatial Data Transfer Standard (SDTS) [12] comprehends not only two-

(one- and zero-) but also three-dimensional universals (called ‘entity types’) such as *fumarole*, *grave*, *mount* and *trough*.

While the FMA deals primarily with material objects and their boundaries, it also deals with portions of body substances (e.g. of water, urine, or menstrual fluid) and with the body spaces (cavities, conduits) which these occupy [13,14]. GIScience deals similarly not only with material geographic objects such as mountains and forests, but also with non-material geographic objects (such as valleys and craters) having some of the features of containers or conduits. In an extended sense it deals also with the substances (above all portions of fresh and salt water) which occupy these.

Considerable progress has been made on the geospatial side, not only in the standardization of geospatial terminology, but also in the development of formal theories and tools for both quantitative and qualitative spatial reasoning [15,16, 17,18,], theories and tools which have since been applied in the biological domain [19, 20]. In the geospatial case, the tools in question are applied primarily in reasoning about the fixed spatial regions on the surface of the earth with which spatial objects can be associated. Analogous tools for region-based reasoning are more difficult to develop and apply in the domain of anatomy because of the elasticity of the human body as contrasted with the earth as base reference object [4]. (It is an essential feature of the heart, for example, that it is constantly in the process of becoming spatially deformed.) On the other side, however, geospatial ontology is less advanced than anatomy in that it has nothing like the formal sophistication in its treatment of ontological relations, and nothing comparable to the coverage, in terms of systematicity and number of universals treated, that is manifested by the FMA. Thus the SDTS comprehends only some 200 entity types.

5 The Proper Treatment of Relations in Ontologies

We can conceive ontologies for present purposes as *controlled, structured vocabularies* designed to support the integration of data and information deriving from heterogeneous sources. An ontology like that of the FMA is structured through assertions of the form ‘*A relation B*’ (where ‘*A*’ and ‘*B*’ are terms in the FMA vocabulary and ‘*relation*’ stands in for ‘*part_of*’ or some similar expression). Such assertions express general statements about the corresponding universals, which correspond to the sorts of statements found in scientific textbooks. To link such ontologies to reality, however, we need to take account not only of the universals described in scientific theories but also of the corresponding instances or tokens which we find about us in reality, and this means that we need to deal not only with the universal–universal relations commonly treated of in work in ontology, but also with instance–universal and instance–instance relations [7].

Thus for example the thesis according to which *lobe of liver part_of liver* – which expresses a universal-universal relation – gets its reference to reality in virtue of the fact that it is in part a thesis about instances, to the effect that:

every canonical instance of the universal *lobe of liver* is a part (in the instance-level, i.e. in the standard mereological, sense of ‘part’) of some instance of the

universal *liver*.

Note the all–some structure of this assertion, which is copied also in parallel universal-universal assertions involving other spatial relations such as adjacency, attachment and continuity [5,21].

As already noted, one important distinction between geospatial and anatomical spatial reasoning turns on the different roles that universals and instances play in their respective map-like representations. On the geospatial side, universals are captured in the *legends* that we find in the corners of maps, legends which can themselves be seen as forming miniature ontologies in their own right. In the anatomical side, in contrast, universals play a central role, not least in virtue of the tremendous increases in our knowledge (for example under the auspices of the Human Genome Project) about the ways in which universals on the coarse-anatomical levels are connected to universals at level of finer grains, down to molecules. Thus where the question of which universals need to be distinguished by an anatomy ontology was once resolved by visual inspection, scientific anatomy rests increasingly on empirical research in the domain of genetics. The parts of the body demarcated on the basis of phenomenology will in the future be acknowledged as genuine anatomical structures only after it has been demonstrated that there are structural genes whose coordinated expression in the development of organisms of the corresponding types brought forth the relevant instances. Hence the FMA in its full version must contain a place also for developmental transformations.

6 Anatomical Entities

Four upper-level universals of the FMA are: *anatomical structure*, *substance*, *space*, and *boundary*.

Anatomical structures are *material* entities such as organisms, organs, cells, and biological macromolecules, which have their own inherent three-dimensional shapes.

Body substances are portions of blood, water, urine or cerebrospinal fluid, entities which inherit their three-dimensional shapes from whatever are the relevant containers. The portion of blood in your right ventricular cavity at some specific time has a shape which it inherits from the surrounding ventricle.

Body spaces are *immaterial* anatomical entities (cavities, orifices, conduits), whose shape, again, is inherited from the relevant surrounding anatomical structure. They are distinguished from spatial regions in that they are *parts of organisms*, which means that they move from one spatial region to another with the movements of their hosts.

Anatomical boundaries are distinguished from anatomical entities in the other three classes by the fact that they are of lower dimension, and stand in a relation of boundary dependence [22] upon some relevant anatomical structure or landmark [6].

7 Anatomical Structures

Mereotopologically speaking, anatomical structures are marked by the fact that they

are maximally self-connected, which means that they have their own complete three-dimensional connected physical or bona fide boundaries. Virtually all anatomical structures, however, are connected to neighboring anatomical structures via conduits which link the anatomical spaces within them. To take account of this fact we require a reading of ‘maximally self-connected’ that allows for corresponding portions of fiat boundaries – we can think of these as punctures in their external surfaces whose areas are quite small in relation to the corresponding total boundary. (‘Fiat boundary’, here, signifies a boundary in a continuant entity which corresponds to no physical discontinuity or bona fide boundary in the entity itself, but rather to a delineation which is drawn by human beings. [10])

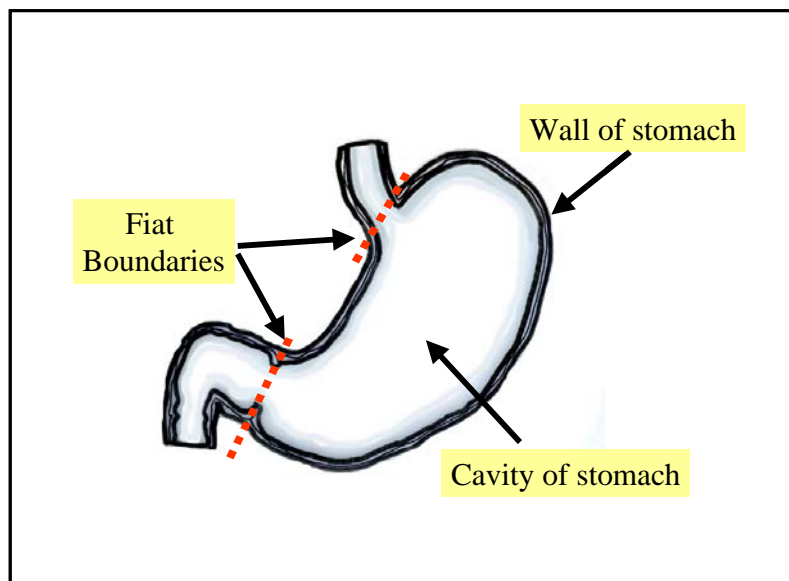


Figure 1. The stomach and its major outlets

Which small portions of fiat boundaries we can ignore in specifying anatomical structures is not a trivial matter. Consider the small portions of fiat boundaries we need to allow in delimiting an anatomical structure such as a stomach or kidney (see Figs. 1 and 2). The stomach, we might think, would be an unproblematic fiat entity, because it is merely a segment of a certain tubular continuum which includes also for example the esophagus and the small and large intestines. In some cases we can find bona fide *landmarks*, specific changes within the mucosal and muscular layers which form the walls of the relevant cavities, for the drawing of the fiat boundaries which extend laterally across the relevant cavity. However, the cases of the stomach and esophagus and of the kidney and ureter show that not all anatomical structures, at those places on their surface where they are demarcated by fiat boundaries, are demarcated by fiat boundaries which are *landmarked* in this sense.

In the case of the kidney we have an anatomical structure that is separated from its surroundings largely by bona fide boundaries and only by small sections of fiat boundary, but in such a way that the fiat boundary in question is located in a non-fiat entity (the urinary tract) (Figure 2).




Figure 2. The kidney with vertical bar designating a fiat boundary in relation to the arterial and venous systems

8 Fiat Boundaries and Partitions

Fiat boundaries come in two types: those which demarcate physical anatomical entities (for example *the plane of the esophagogastric junction*, which demarcates the esophagus from the stomach) and those which demarcate anatomical spaces (for example *the plane of the pelvic inlet*, which demarcates the abdominal from the pelvic cavity). Our talk of '*planes*,' here, draws attention to the fact that most anatomical fiat boundaries have geometrically regular shapes, just as is often the case in the geospatial realm (the borders of Colorado or Manitoba).

In addition to anatomical structures in the technical sense of the FMA, we can recognize also:

1. fiat *parts* of anatomical structures (for example the fundus of the stomach), which are not complete;
2. fiat *aggregates* of anatomical structures (for example the aggregate of the upper and lower limbs), which are not connected themselves.

The recognition of fiat entities of these sorts allows us to do justice to the fact that one

and the same anatomical structure can be partitioned in different ways [23]. The stomach can be decomposed in one context into its fundus, body and pyloric antrum and in another into its wall and cavity. The FMA sees the former as a fiat partition into *regional parts*, the latter as what it calls a ‘compositional partition’ into *constitutional parts* [24]. While constitutional parts are genetically determined, regional parts – for example the loin or the epigastrium – are defined in part by arbitrary coordinates. Even if we remain with bona fide parts, however, then we need to acknowledge cross-cutting demarcations. Thus for example the bona fide boundary between the gray and the white matter of the brain is cross-cut by the bona fide boundaries of the neurons which pass between them.

9 Connectedness and Continuity

The body’s component parts are intimately interconnected. Indeed, if we leave aside the cells floating free in blood and other body substances, then practically all anatomical structures are connected to other anatomical structures through different kinds of continuities or junctions.

The FMA analyzes the relation of connectedness in terms of three different kinds of relations: *continuous_with*, *attached_to* (e.g. of muscle to bone) and *synapsed_with* (of nerve to nerve and nerve to muscle – a special type of attachment relation obtaining at the level of granularity of cell parts, axons and dendrites).

Two continuants are *continuous* on the instance level if and only if they share a fiat boundary. A continuant is *self-connected* if and only if any division of the entity yields parts which are continuous. The relation of continuity on the instance level is of course always symmetric. On the class level, however, this is not the case. To see why not, consider the relation between the lymph node and the lymphatic vessel. Each lymph node is continuous with some lymphatic vessel, but there are lymphatic vessels (e.g. lymphatic trunks such as the thoracic duct) which do not stand in continuous connection to any lymph nodes. We thus have *lymph node continuous_with lymphatic vessel* (because for every instance of the former there is some instance of the latter with which the former is continuous), but not *lymphatic vessel continuous_with lymph node*.

To understand the relation *attached_to*, consider the junction depicted macroscopically in Figure 3, which shows a bone and a muscle, the latter consisting of a tendon and a muscle belly, and (on a finer-grained level) of collagen fibers, muscle fibers and bone matrix. The bone itself is well delimited: it ends where the bone matrix ends. The same applies to the muscle fibers which, due to their contractile elements, are clearly demarcated from the tendon. But collagen fibers cross all of these boundaries. One fiber might overlap with the muscle fascia and the tendon, another with the tendon and the bone. Tendon and bone can be separated only by severing the fibers in question.

Attachment, too, is symmetrical on the instance level. On the class level, however, it is not in general a symmetrical connectivity relation. To see why not, consider the universals *placenta* and *uterus*. Every instance of the former is indeed attached to some instance of the latter; not, however, conversely.

While the corresponding instances have their own bona fide boundaries, the distal tendon comes into intimate contact with that circumscribed area of the bone where extensions of its collagen fiber bundles of the tendon (the Sharpey's fibers in Figure 3) penetrate the bone and intermingle with collagen fibers in the bone's own matrix. The tendon may thus be separated from the bone only by severing Sharpey's fibers.

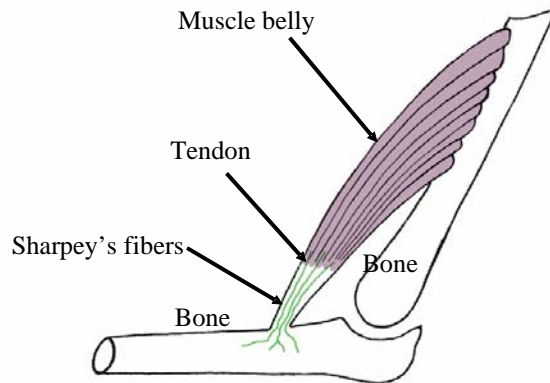


Figure 3. An anatomical junction

10 Location and Containment

In addition to the relations of instantiation (between an instance and a universal) and parthood (between one continuant instance and another continuant instance), the FMA contains also a treatment of location. To understand location formally, we associate with the human body a collection of regions (*relative places* in Donnelly's terms [25]) and define a function which assigns to each anatomical entity c and time t the corresponding region $r(c, t)$ which c exactly occupies at t . We can then define the relation of location for anatomical instances as follows: [26]

c **located_in** d **at** t =def. $r(c, t)$ **part_of** $r(d, t)$ **at** t

where relations picked out in **bold** obtain at the instance level. On the level of universals we have:

C **located_in** D =def for all c, t , if c **instance_of** C **at** t then there is some d such that d **instance_of** D **at** t and c **located_in** d **at** t .

Trivially, by this definition, all parts of anatomical structures are located in the corresponding wholes. But there is a second type of location relation distinguished in the FMA, which is that of *containment*. This holds between material anatomical entities (body substances and anatomical structures) and the anatomical spaces in which they are contained. For example: the right lung is contained in the right half of

the thoracic cavity; a Ca^{++} ion is contained in an intracellular space in a heart muscle cell.

If we move away from canonical anatomy, we encounter cases where, for example, a lobe of a liver is removed from a donor and transplanted into a hepatectomized patient. There is then an instance of the universal *lobe of liver* which is for a certain time not a part of any instance of *liver*. However, even when such non-canonical cases are taken into account, it still has to be true that every instance of *lobe of liver* stands in an instance-level parthood relation to some instance of *liver* at some time (more precisely still: every instance of *lobe of liver* stands at the beginning of its existence in such a relation to some instance of *liver*).

Two entities may thus be related in terms of parthood only during a certain phase of their simultaneous existence. It is for this reason that parthood relations between continuants on the instance level must be indexed by times [7]. This is not specific to living systems. For example, a screw can be part of an engine and it can then be substituted by another, replacement screw. In contrast to artifacts, however, biological objects are engaged in a constant exchange of matter with their environment, so that many parthood relationships are short-lived. Moreover, the dynamic phenomena of matter exchange [27] indicate that there must be relations intermediate between parthood and containment, realized for example when you take a bite out of an apple (and the relevant portion of apple moves from being part of the apple to being contained in your oral cavity). Consider what happens when food is degraded in course of digestion into sugars, amino acids and fatty acids. Those portions of such substances in the lumen of the stomach are then *contained* therein, while those that have traversed the epithelium are successively *parts* of epithelial cells and then of blood or lymph.

11 Criteria of Parthood

An account of the processes in question must accordingly allow for the existence of transitions between containment and parthood. How, given the above, are we to distinguish genuine parthood from the relation of being merely spatially included within (i.e. from the relation **located_in** as defined above)? Is an embryo part of, or merely located in, a uterus? [28] Is a bolus of food part of, or merely located in, a digestive tract? Is an oxygen molecule part of, or merely located in, a lung? We here offer four kinds of criteria which may be of assistance in answering such questions (for further detail see [29]).

1. *Genetics*: The parts of the body should be of the same genetic origin as the body itself. Thus the embryo, on this criterion, is not a part of the body of the mother. This criterion faces problems for example in application to oxygen or nitrogen molecules in the body (since these do not have a genetic origin) or to the mitochondria found in nearly every cell of the body (which have their own DNA).
2. *Sortality*: If continuant *c* is part of continuant *d*, then *c* and *d* must be of the right sorts to make this possible (they must instantiate appropriate universals). Thus if *d* is an organism, then it is ruled out that *a* should be an artifact (e.g., a

heart pacemaker, a bullet), or a second whole organism (a symbiont, parasite, prey, embryo or fetus).

3. *Life Cycle*: Unless this is already ruled out by sortal constraints, we can infer from the fact that c is located in d during the whole of the life cycle of d , that c is part of d (the right ventricle of the heart is for this reason part of the heart).

4. *Function*: We can infer parthood from location, finally, where an object c , located in a second object d , has a function whose exercise or performance is essential to d 's survival or to the maintenance of d 's proper functioning. The functioning of the heart or of the brain is essential for the survival of the whole human body in this sense, where a given volume of urine is not essential to the survival of the bladder, and hence the urine is not counted as part of the bladder. This criterion faces problems for example in application to hair (is hair essential to the proper functioning of the human body?) or to the kidney and of those other organs supplied to the body in pairs [30].

12 Holes and Parts

A further family of problems connected with the location relation in anatomy turns on the fact that the boundaries of the objects with which we have to deal may themselves be difficult to specify. Many anatomical objects are sponge-like; they are replete with vessels, capillaries, cavities, holes and ducts of various sorts. This is true of all the body apart from the cornea and lens. A clear delimitation of an anatomical object, often including reference to a plurality of distinct levels of granularity, is therefore essential for making any assertion about location.

Is a small object such as a calculus located in a duct inside a gland also located in the gland itself, or, in the case when the lumen of the duct communicates directly with the exterior, located only in the exterior surrounding space? The answer to this question depends on whether we admit spaces as parts of anatomical objects in cases such as this.

The range of problematic borderline cases connected with surface structures is depicted in Figure 4. Of the white and gray volumes falling below the (rough) line of demarcation of the surface of the body in question, which are parts of the exterior of this body and which are parts of the body itself?

Analogous puzzles arise also in connection with spatial discontinuities. Accessory spleens such as are illustrated in stylized fashion in Figure 5 can be found in more than 10% of the population. This phenomenon can be accounted for in two ways, either by admitting one discontinuous object with three parts, or (with the FMA – which is generally averse to the admission of discontinuous anatomical entities) by admitting three distinct objects which collaborate in the exercise of a certain function.




Figure 4: Problematic surface structures

In sum, the specification of anatomical part, location and connection relations, and also of the degree of spatial overlap between anatomical structures, is often problematic because the relevant spatial extensions are difficult to delimit and because the relevant anatomical entities continuously lose and gain parts and continuously exchange matter with their environment. It will likely never be the case that we can formulate a criterion for parthood that can be guaranteed to yield a determinate answer to the question *is x a part of y?* in every single case. From this some might be tempted to conclude that the notion of parthood, at least in the biological domain, has



Figure 5: Normal spleen (left); accessory spleens (right)

an ineliminable element of indeterminacy. We, however, prefer to see the indeterminacy as lying rather in our partial knowledge of the relations between the corresponding entities in reality.

13 Non-Canonical Anatomy

The FMA is a representation of canonical anatomy, which means that its individual variables range over those adult male and female human beings who satisfy the generalizations which appear in textbooks of structural anatomy and which conform to a pattern repeatedly observed by many generations of anatomists and surgeons over several centuries. By appealing to the device of specifying different ranges of variables, we can modify the scope of the FMA to represent generalizations belonging to the different branches of anatomy, for example to canonical human beings at various stages of embryological development, and even to organisms of other species. It can allow us also to represent the generalizations governing the anatomical variants yielded by the presence of, for example, coronary arteries or bronchopulmonary segments, which deviate from canonical anatomical patterns of organization in various well-understood ways. We here conclude with some brief and speculative remarks on the proper treatment of pathological structures within an ontology like the FMA.

While the universal *colon* as it appears in the FMA comprehends only instances of normal colons (however the term ‘normal’ is to be defined), if the FMA ontology is to serve as a basis also for *non-canonical* anatomy then it must have the facility to extend this range of instances to include also abnormal colons. The resulting framework of *pathological anatomy* might then include assertions such as

abnormal colon is_a colon;

colon carcinoma pathological structure part_of abnormal colon.

Moreover (recalling our ‘all–some’ reading of the class-level relation *part_of* above), we might have:

colon carcinoma pathological structure part_of colon.

Every colon carcinoma is part of some colon, even though not every colon has some colon carcinoma as part.

We might have

abnormal colon transformation_of colon

colon with carcinoma is_a abnormal colon

colon with carcinoma transformation_of normal colon

where *C transformation of D* is defined as obtaining whenever *C* and *D* are continuant universals which are such that every instance of *C* was also an instance of *D* at some earlier time in its existence [7].

14 Towards Anatomical Information Science

Leaving aside a number of abstract domain-independent treatments of spatial structures and relations, the primary focus of the discipline called ‘Spatial Information Theory’ has thus far been in the area of geospatial information. We believe however that Spatial Information Theory ought also to encompass the theory of spatial

properties and relations in other domains, not least because – as we hope to have shown in the foregoing – the latter can introduce important phenomena of a kind not thus far considered in the literature.

Because of the special role of the canonical (idealized) human body, and because of the complementary special role of variant and pathological anatomic structures in anatomical information science, many features of the type of spatial information science realized in the FMA will be unfamiliar to those working on spatial representation and reasoning in the geographical domain. In cartographic terms, canonical anatomy would correspond to a map of an idealized portion of geographic reality (an idealized city, say, or an idealized lake or continent). Corresponding to actual maps (of actual cities or actual continents) is *instantiated anatomy*, which comprehends anatomical data about actual human beings of a type that might be recorded in a clinical record or captured in a radiographic image [9]. Instantiated anatomy deals with individual, living, human subjects, but in a way that relies on the categories or kinds depicted in canonical anatomy. Practically all of geography is instantiated geography and geospatial information science is in consequence characterized by the existence of a large mass of spatially referenced instance data and of powerful systems for reasoning with this data, combined with a treatment of the corresponding universals which is relatively impoverished from the theoretical point of view.

There is also a *normative* dimension of the discipline of canonical anatomy, which has no direct counterpart in the geospatial domain. For while there are healthy and unhealthy cities, it is not the case that all healthy cities have a more or less identical groundplan. Moreover, the geospatial domain has no counterpart of the contemporary evidence-based discipline of medicine, and thus no counterpart of its central organizing discipline of canonical anatomical science (and thus no scientific interest in, for example, maps of ideal cities). But for this reason, too, there is no counterpart of pathological anatomy in the domain of geospatial science, which is to say: no science of the determinate ways in which geospatial entities such as cities or lakes depart from some normative (‘normal’) case.

Thus the SDTS contains within its list of attributes no terms for what we might think of as disorders of its entity types, as contrasted with the 900,000 or so terms included in SNOMED-CT, the systematized nomenclature of clinical terms maintained by the College of American Pathologists [31], a large fraction of which refers to disorders. There is no counterpart, either, of the ways in which human anatomy can be related to the anatomy of other species as a basis for the detection of what may be medically relevant homologies [32].

The existence of the FMA means that anatomical information science rests on an impressive tool for the treatment of anatomical universals, even though both the associated instance data and the tools for reasoning with such instance data are still impoverished. Some progress is being made on the side of instance-level anatomical information science. Again, however, the problems of elasticity, movement, and growth of bodily organs present considerable obstacles to the development of corresponding tools for instance-based spatial reasoning [33]. With the development of genomics-based individualized medicine, and with associated increases in the sophistication of electronic health records and medical image analysis, we believe that the imbalance between class- and instance-based anatomical data will in the coming

years be gradually resolved. Corresponding tools for representation and reasoning with anatomical instance data will thus increasingly be needed [34], and we can anticipate that the FMA will play an important role in their development by being used in tandem with some of the reasoning tools developed in recent years in the spatial domain.

Acknowledgments: This paper was written under the auspices of the Wolfgang Paul Program of the Alexander von Humboldt Foundation, the European Union Network of Excellence on Semantic Interoperability and Data Mining in Biomedicine and the Volkswagen Foundation under the auspices of the project “Forms of Life.”

References

1. <http://sig.biostr.washington.edu/projects/fm/>
2. Rosse, C. and Mejino, J. L. V. Jr. A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy. *J Biomed Informatics*, 2003; 36:478–500.
3. Noy, N. F., Mejino, J. L. V. Jr., Musen, M. A. and Rosse, C. Pushing the Envelope: Challenges in Frame-Based Representation of Human Anatomy. *Data and Knowledge Engineering* 48 (2004) 335–359.
4. Grenon, P. and Smith, B. SNAP and SPAN: Towards Dynamic Spatial Ontology. *Spatial Cognition and Computation*, 4: 1 (March 2004), 69–103.
5. Cook, D. L., Mejino, J. L. V. Jr. and Rosse, C. Evolution of a Foundational Model of Physiology: Symbolic Representation for Functional Bioinformatics. *Proceedings MedInfo* 2004: 336–340.
6. Mejino, J. L. V. and Rosse C. Symbolic modeling of structural relationships in the Foundational Model of Anatomy, *Proceedings of KR-MED 2004* (First International Workshop on Formal Biomedical Knowledge Representation), 48–62.
7. Smith, B., Ceusters, W., Klagges, B., Köhler, J., Kumar, A., Lomax, J., Mungall, C., Neuhaus, F., Rector, A., Rosse, C. Relations in Biomedical Ontologies. *Genome Biology*, 2005, 6 (5), R46.
8. Donnelly, M., Bittner, T., and Rosse, C. A Formal Theory for Spatial Representation and Reasoning in Biomedical Ontologies. *IFOMIS Reprints* 2005.
9. Rosse, C., Mejino, J. L., Modayur, B. R., Jakobovits, R., Hinshaw, K. P., Brinkley, J. F. Motivation and Organizational Principles for Anatomical Knowledge Representation: The Digital Anatomist Symbolic Knowledge Base. *Journal of the American Medical Informatics Association* 1998; 5: 17–40.
10. Smith, B. On Drawing Lines on a Map, in Andrew U. Frank and Werner Kuhn (eds.), *Spatial Information Theory. A Theoretical Basis for GIS* (Lecture Notes in Computer Science 988), Berlin/Heidelberg/New York: Springer, 1995, 475–484.
11. Smith, B. and Varzi, A. C. Fiat and Bona Fide Boundaries, *Philosophy and Phenomenological Research*, 60: 2, March 2000, 401–420.
12. US Geological Survey Spatial Data Transfer Standard (SDTS) Information Site (<http://mcmcweb.er.usgs.gov/sdts/>).
13. Casati, R. and Varzi, A. C. *Holes and Other Superficialities*, Cambridge, MA: MIT Press 1994.

14. Donnelly, M. On Parts and Holes: The Spatial Structure of the Human Body, in: M. Fieschi, E. Coiera, and Y. J. Li (eds.) *Proceedings Medinfo 2004*, 351–356.
15. Cohn, A. G., Bennett, B., Gooday, J. and Gotts, N. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus, *Geoinformatica 1*, 1997, 1–44.
16. Smith, B. Mereotopology: A Theory of Parts and Boundaries, *Data and Knowledge Engineering* 1996; 20: 287–303.
17. Donnelly, M. Layered Mereotopology. *Proc IJCAI 2003*, 1269–1274.
18. Casati, R. and Varzi, A. C. *Parts and Places: The Structures of Spatial Representation*, MIT Press, Cambridge, 1999.
19. Bittner, T. Axioms for Parthood and Containment Relations in Bio-Ontologies. *Proceedings of KR-MED 2004* (First International Workshop on Formal Biomedical Knowledge Representation), 2004; 4–11.
20. Schulz, S. and Hahn, U. Mereotopological Reasoning about Parts and (W)holes in Bio-Ontologies, *Proceedings of FOIS*, New York: ACM Press, 2001: 198–209.
21. Smith, B. and Rosse, C. The Role of Foundational Relations in the Alignment of Biomedical Ontologies, *Proceedings of Medinfo 2004*; 444–448.
22. Smith, B. On Substances, Accidents and Universals. In Defence of a Constituent Ontology, *Philosophical Papers*, 27 (1997), 105–127.
23. Bittner, T. and Smith, B. A Theory of Granular Partitions, *Foundations of Geographic Information Science*, London: Taylor & Francis, 2003.
24. Mejino, J. L. V., Agoncillo, A.V., Rickard K. L. and Rosse C. Representing Complexity in Part-Whole Relationships within the Foundational Model of Anatomy, *Proceedings of AMIA Symp.* 2003; 450–454.
25. Donnelly, M. Relative Places. In A. Varzi and L. Vieu (eds.) *Formal Ontology in Information Systems, Proceedings of FOIS 2004*, Amsterdam: IOS Press, 249–260.
26. Donnelly, M. A Formal Theory for Reasoning about Parthood, Connection, and Location. *Artificial Intelligence* 160 (2004) 145–172.
27. Cohn, A. G. Formalizing Bio-Spatial Knowledge, *Proceedings of FOIS 2001*, ACM Press, New York, 2001, 198–209.
28. Smith, B. and Brogaard, B. Sixteen Days, *The Journal of Medicine and Philosophy*, 28 (2003), 45–78.
29. Schulz, S., Daumke, P., Smith, B., Hahn, U. How to Distinguish Parthood from Location in Bioontologies, *Proceedings of AMIA Symposium*, 2005, in press.
30. Johansson, I., Smith, B., Munn, K., Tsikolia, N., Elsner, K., Ernst, D. and Siebert, D. Functional Anatomy: A Taxonomic Proposal, *Acta Biotheoretica*, in press.
31. <http://www.snomed.org>
32. Travillian, R. S., Rosse, C. and Shapiro, L. G. An Approach to the Anatomical Correlation of Species through the Foundational Model of Anatomy. *Proceedings of AMIA Symposium*, 2003; 669–673.
33. Pilgram, R., Fritscher, K. D., Fletcher, P. T., and Schubert, R. Shape Modelling of the Multiobject Organ Heart, *IATED: International Conference on Biomedical Engineering – BioMED 2004*, Acta Press, 2004, 157–160.
34. Ceusters, W. and Smith, B. Tracking Referents in Electronic Healthcare Records, *Proceedings of Medical Informatics Europe*, 2005, in press.