SEMANTIC RELATIONS IN A MEDICAL DIGITAL LIBRARY

Nina Wacholder*, Celina Imielinska**, Judith Klavans*, Ewa Soliz*, Pat Molholt*
*Columbia University   **Stevens Institute of Technology

Vesalius™ Project
Office of Scholarly Resources
Columbia University Health Sciences
701 W. 168 St., Room 201
New York, NY 10032
nina@cs.columbia.edu   celina@menger.eecs.stevens-tech.edu

1 Abstract

In this paper, we describe the Vesalius™ Project, a multi-modal collection of anatomical resources under development at Columbia University. [1] Our focus is on the need for navigational tools to effectively access the wealth of electronic information on anatomy, including life-like 3D images of anatomical entities that can be interactively viewed and browsed. We describe a key component which must be in place in order to develop a flexible and reusable digital library system, namely an anatomical knowledge base containing a 'nucleus' of anatomical information specifically designed to make it possible to develop a wide spectrum of curriculum applications that use and extend the information in the knowledge base. The unique contribution of our research lies in the dual focus on user needs and on effective use of knowledge representation theory in order to develop a system that makes it possible to take advantage of interactive 3D models and the wealth of other anatomical data now available.

2 3D images for medical education

In this paper, we address the problem of how to organize 3D anatomical images and related information in a digital library intended specifically for use in a medical education environment. [2] The impetus for this work is the availability of realistic 3D models of anatomical entities of sufficiently high quality that medical school faculty are willing to use them in the classroom, and of sufficient portability that they can be examined interactively on computer platforms available to graduate students. The most important of these visualizations are 'true' 3D models that we have reconstructed from the Visible Human dataset of high resolution digitized cadavers [1], [2], [3], [4], [5]. A snapshot showing a single view of the heart reconstructed from this dataset is shown in Fig. 1.

The ability to browse the human anatomy in three dimensional space is especially important for first year medical students, who are under pressure to quickly develop a mental picture of the human body, with all of its spatial relationships. This task is notoriously difficult with 2-D images alone, because students must mentally translate them into 3-D ones. The dissection lab, while important, also does not suffice: the tissue color and texture of cadavers is quite different than that of living beings; many delicate structures are destroyed by the process of dissection; and finally, the cadaver, once taken apart, can not readily be put back together again.

The new techniques for displaying these 3D images that we discuss in this paper give us the ability to view anatomical entities in relation to other entities and their subparts. However, our experience with these images has shown us that images by them-

---

[1] Trademark held by Columbia University. The Vesalius™ Project is named after Andreas Vesalius, a sixteenth century anatomist whose work laid the foundations for all subsequent anatomical research. The Columbia University Health Sciences Library owns several first editions of his work.

[2] Stephen Johnson, a member of the Vesalius™ team, made helpful comments on an earlier version of this paper and

---

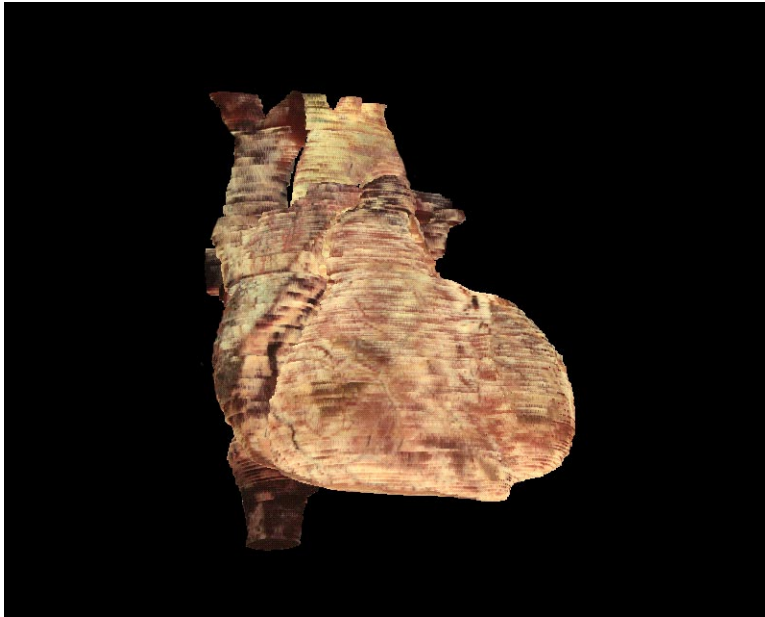Jon Morrow assisted with verification of medical terms and concepts.

Figure 1: Lateral view of the heart. Vesalius^{TM}

selves are not enough in the medical education context. The Vesalius^{TM} Project, currently under way at the Columbia University College of Physicians & Surgeons' Office of Scholarly Resources (P&S), is in the process of building a digital library of multi-modal resources organized around the 3D images. In addition to the anatomical images themselves, the resources that we are collecting for this multi-modal system include video and audio material, 2D slides and photos, explanatory text, lectures and test material [6], [7], [8].

The research we describe in this paper is based upon the hypothesis that a semantic network of anatomical information is a necessary preliminary to the development of a navigational system which reflects a conceputally sound organization of the material and is also computationally tractable. Specifically, we discuss the kinds of navigation that our users need and we discuss some of the properties of an anatomical ontology designed to make it possible for users to navigate part-whole relations, function relations and other meaningful semantic relations. The semantic network will reside in a knowledge base specifically designed to support a wide spectrum of curriculum applications that use, extend and/or filter this information.

In the next section, we describe the types of 3D images that will reside in the knowledge base. Then we discuss issues that arise in navigating these images; finally we turn to the semantic network that will linked to the image base and the text base.

3   Obtaining 3D anatomical images

An important goal of the Vesalius^{TM} Project, a collaborative effort between Columbia University and the Stevens Institute of Technology, is to refine the available technology in order to obtain the highest possible photographic-quality smooth models of anatomy which can be manipulated in real time on any platform and transmitted efficiently [3], [4], [5]. The Visible Human dataset consists of high resolution color digitized slices of frozen female and male cadavers, where the anatomical detail of each slice is 1/3x1/3 mm. [1] [2]. Because of the quality of the dataset, the 3D visualizations of anatomical structures can capture both 'actual' color and fine detail.

The 3D models of anatomical structures are obtained from the dataset by a complex process. We start with color segmentation of the target anatomical structure in each individual slice, through the range of slices that enclose the structure [9]. We then reconstruct the 3D surface using the extended color alligator algorithm [10] [11], where vertices of every triangle on the surface are 'dipped' in the original color. Using a virtual camera model, 3D anatomical structures are

shaded and projected on a 2D viewing screen. Because we preserve the color of the original dataset, the 3D models look like photographs of the actual anatomical structures 'carved' from the frozen body. These models can be viewed from arbitrary angles. Fig. 2 shows 36 snapshots of the heart, each from a different perspective.

The quality of the 3D models reconstructed from the Visible Human dataset reflects the state-of-the-art of existing visualization tools and the actual level of detail (LOD) of the dataset. Therefore, some of the visualizations of very fine structures may not satisfy anatomists' expectations. For this reason, we are supplementing the true 3D models with other highly photographic visual material. For example, we have developed an application using interactive 2D photos of the skull, and we are building 3D objects from 2D photographs of a physical skull. The trade-off is between fully interactive 3D models that are not completely smooth or truly photographic *or* truly photographic 3D images which are not fully interactive 3D models.

From the point of view of anatomy, an important property of the visual material created by the techniques described above is that they can be used to represent visible anatomical structures and their subparts in relation to each other. Any subsets of the collection of 3D objects reconstructed from the Visible Human dataset can be presented individually or in groups (scenes) and viewed from arbitrary perspectives. For example, the aorta can be viewed as part of the heart, as in Figs. 1 and 2, or it can be viewed in isolation, as in Fig. 3 where part of the aorta has been 'plucked out' of the heart for inspection. Similarly, the photographic skull can be separated into its component bones and each bone can be viewed and rotated individually.

These visualization techniques give us the ability to present subparts in relationship to each other or to a larger whole. Our long-term goal is the reconstruction of all anatomical structures together with their medically interesting subparts. For now, we are developing prototypes of the heart and the skull.

4   Navigating anatomical information

In working with 3D anatomical images, we have found that there is an urgent need for metadata to helps users navigate the images, and to link visual material with textual material. The three-dimensional quality of the images adds to the possible number of views that the system can make available. The user needs to be able to specify which view of an organ is of interest, e.g., a lateral or frontal view. Furthermore, the ability to display meaningful subparts of anatomical entities in relation to each other as well as to the entities of which they are part brings with it a host of special navigational issues. For example, the aorta extends from the heart, and a user might want to view the aorta in situ or in isolation. To give another example, both the aortic valve (the opening of the aorta) and the cells of which the aorta is made can be said to be part of the heart, the aorta and the rest of the body. However, someone interested in clinical anatomy who is interested in parts of the heart probably wants to look at parts at the aorta, but not at parts of a cell, whereas someone studying histology would want to look at inidividual cells of cardiac muscle or of aortic endothelium, but not at the macroscopic organs themselves.

In our work with users, we have identified a set of properties that are important in navigating a multi-modal digital library of anatomy:

- **Relation between image and text**. The user must be able to use words or text to specify the desired view of the entity, and the system must be able either to take advantage of contextual information in order to decide what view to show (e.g., anterior or posterior) or provide a list of sensible options to the user.

- **Adjacency**. The user may want to first see a structure like the aorta in relation to nearby structures such as the first rib and clavicle which pass near the aorta. Linguistic expressions that convey spatial relations such as *near* or *over* are notoriously ambiguous (e.g., [12]) and hard to formalize for computational purposes, even in two dimensions. The addition of the third dimension further magnifies the problem by increasing the number of possible views denoted by the relationship. The knowledge base must provide a list of adjacent structures to make it easy to access those structures.

- **System vs. region**. In anatomy, the distinction between system and region is basic. When

Figure 2: 36 views of the heart

studying the thoracic vasculature, the user needs to go from the subclavian artery to the maxillary artery; when considering the topology of the thorax, the user should be able to go from the subclavian artery to the first rib, the clavicle, or the anterior scalene muscle.

- **Level of granularity**. In different contexts, different levels of granularity of image and text, are relevant. If a user is interested in the liver, a decision must be made as to the level of granularity of image and text to display, e.g., the macro level where the focus is on physical connections to adjacent entities, or a relatively microscopic level, where the focus is on major and minor calyces, cortex, medulla and fibrous capsule .

- **Attributes**. The symbolic data base must provide a way for the user to compare and contrast structures that differ with respect to many attributes, including developmental stage, structural features, and anatomical abnormalities. Each image must therefore be associated not only with its anatomical components, but also with distinctive anatomical and physiological attributes (i.e., the properties that distinguish one view of a structure from another).

- **Interesting vs. uninteresting information**. The organization must take into account the fact that not all possible views of an anatomical entity are of equal interest. For instance, the distinction between male and female pelvic bones is important, but for cranial bones, it is not. Different information is also relevant at different points in medical education. During the first year, a list of structures of the thorax which are palpable is important; by the second year, that list is too elementary to be of interest.

5 System organization

In order to meet our dual goals of usability and tractability, we have adopted the structure of the human anatomy as the basis for organizing the Vesalius^{TM} system. Using anatomy as the starting point for organization of the knowledge base is logical
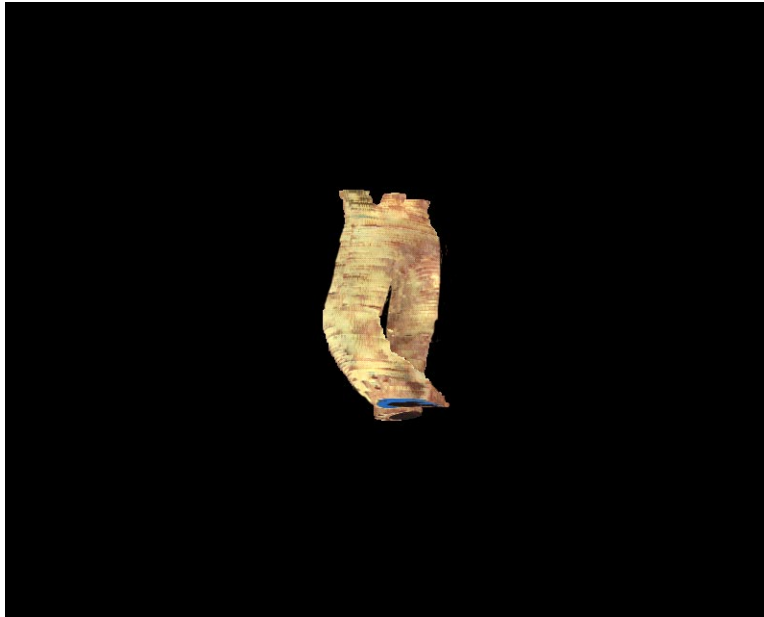
Figure 3: Lateral view of (part of) the aorta. Vesalius$^{\text{TM}}$

because anatomy provides the framework into which medical students traditionally fit the medical knowledge they acquire over the course of their training and careers. First year students study normal anatomy, second year students study abnormal anatomy, which is implicitly and/or explicitly compared to normal anatomy, and third and fourth year students learn to diagnose and treat illnesses by relying on their knowledge of normal and pathological anatomy.

To achieve this integration of anatomical information, the core of the Vesalius$^{\text{TM}}$ system will be implemented as a 'knowledge base', an application-independent mechanism for organizing and linking information. The knowledge base will provide the underpinning for a variety of educational applications such as electronic lecture resources for faculty and independent-study materials for students.

The system will include three components:

1. A knowledge base consisting of information about anatomical entities. The foundation of the knowledge base is a collection of unique medical concepts (identifiers). These will be primarily concepts from anatomy (e.g., chest, heart, pericardium, aorta), with some related information associated with anatomical sites, such as diseases (congestive heart failure, calcification of the peri-

cardium) and procedures (coronary artery bypass). These concepts will be linked in a semantic ontology, a notation for representing information about related concepts in a tree-like graph; the branches of the graph provide links representing meaningful relations between concepts [13], [14].

2. Links to 3D images and other multi-media resources. Each medical concept in the knowledge base can have multiple links to various types of multimedia resources. For example, the concept "heart" can have links to 3D images, 2D models, sections of the online curriculum, and chapters of an online medical textbook. It can also have links to other knowledge-representation systems such as the National Library of Medicine's Unified Medical Language System (UMLS). In this context, a 'link' is simply a string of characters that identifies the location of the multimedia item. For example, if the online curriculum is implemented on the World Wide Web, this location can be the universal resource locator (URL) of the appropriate section. Images and other materials may also be referenced in this manner. An image stored in an image database could be referenced using the identifier (key) employed by that database to store the image.

3. An interface allowing educational applications running on different platforms to access the knowledge base, and to use the links to navigate text and multimedia resources.

6   An anatomical semantic network

In this section, we discuss briefly two important medical knowlege bases which contain anatomical information. Space precludes us from discussing other work such as [15].

The UMLS is the largest and best-known medical knowledge base [16]. The UMLS is an amalgam of several vocabularies developed primarily for literature searching and for use in a clinical environment. The UMLS is an important tool designed to link the literature searching environment with the clinical environment. Although the UMLS integrates several nomenclatures into a single system, at the present time the number of anatomical concepts and the properties of the links are insufficient and the links among these concepts are not defined sufficiently specifically for our use in an educational context.

Important work on "an anatomically based software framework for organizing, visualizing and utilizing biomedical information" is under way at the University of Washington's Digital Anatomist Program [17], [18]. The Digital Anatomist Program, like the Vesalius^TM Project, has the goal of developing an anatomical knowledge base which is flexible and re-usable. [17] has shown that a knowledge-based approach is useful for organization of information about medical images. One of the major applications developed at the University of Washington is the Anatomy Information System, an interactive atlas of anatomy organized by body region (http://www9.biostr.washington.edu/da/html). In this system, the range of anatomical concepts in the UMLS has been expanded and the granularity of the ontological relations among these concepts has been refined. The Digital Anatomist project has been very successful in making important material available but has not concentrated on the problem that we deal with, namely development of ways to navigate images and text. In fact, according to [17]:176, users of the Digital Anatomists' anatomical atlas view it favorably. In addition to more material, the two other major requests have been for "other ways to navigate through the information, and for links to other parts of the anatomy curriculum." A recent article by the Digital Anatomist group takes an approach to developing an ontology similar to ours [19] with careful attention to the classification of anatomical concepts. We have been in contact about the possibility of utilizing each other's results.

To sum, our research shows that available knowledge bases lack adequate coverage of anatomical concepts and rely on relationships that are not formally defined for use with 3D visualizations. In what follows, we discuss some of the properties of semantic relationships that are needed for an anatomical ontology that can be used as the basis for a system for navigating anatomical information.

6.1   The is-a relation

The *is-a* relation is the most important relation in both the UMLS and the Digital Anatomist system. It is well-known that the linguistic expression 'is a' covers a variety of semantic relations (for a classic discussion of these relations see [20]). Semantic networks that use *is-a* relationships usually place a formal restriction on the *is-a* relationship by defining it not in linguistic terms but in terms of inheritance: in a semantic network, the subtype inherits all properties of the supertypes and siblings have distinguishing features.

The *is-a* relationship with the formal restriction on inheritance is very useful for taxonomical (classificatory) purposes. It is widely used because its properties are have been extensively studied and are therefore relatively predictable. The *is-a* relation is particularly suited for the representation of relationships that are strictly hierarchical. For example, the structure of the blood system can be represented relatively naturally, as shown in Fig. 4 [18].

However, a semantic network populated primarily by *is-a* relations, such as the UMLS, does not provide sufficiently flexible navigation from one type of anatomical entity to another. For example, the user might want to go from a vein to a nearby nerve; the hierarchical network does not make it easy to move directly from system to system or region to region.

Moreover, the UMLS hierarchy is not designed in a way which makes it easy to compare information about normal or canonical anatomical entities with information about abnormal ones. At the highest level, the UMLS' hierarchy of semantic types for anatomy
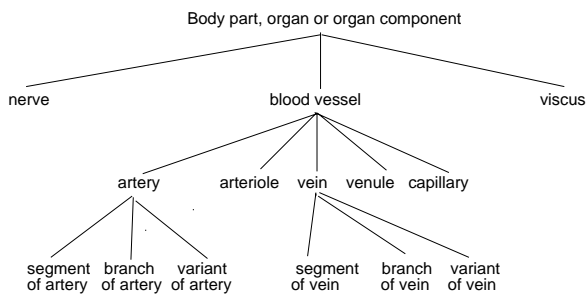
Figure 4: University of Washington's semantic types for thoracic vasculature

is shown in Figure 5, adapted from [21]. With a tree like this, a normal heart and a heart with a calcified pericardium are not linked, except at the highest level in that they are both a type of heart. The semantic ontology must be modified in order to make it easier to establish links between normal and abnormal anatomical entities.
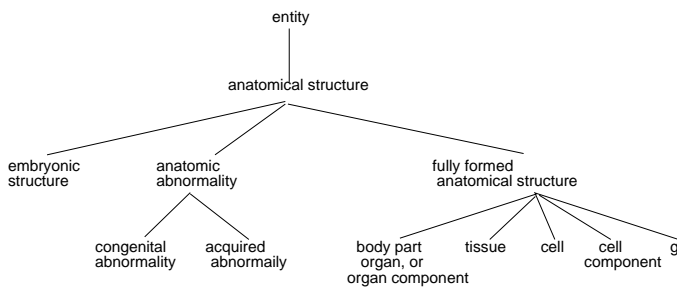

Figure 5: UMLS hierarchy of anatomical entities (top-level)

6.2 The part-of relation

In addition to classification, one important need our users will have is to move among regions and systems. One way to provide this flexibility is with *part-of* or meronymic relations. Like 'is a', the linguistic expression 'part of' (e,g. Winston et al. [22] is multiply ambiguous. Among the types of *part-of* relations that are relevant to anatomy are the following:

- component-integral object (components form patterned structure)

  The heart is part of the body.

- stuff-object (what something is made of)

  The heart is made of muscle tissue.

- place-area (part of mass, but can't be separated)

  The apex of the heart is part of the heart.

In practice, medical networks define the *part-of relation* loosely or not at all. For example, the UMLS defines the *part-of* relation as "composes, with one or more physical units, some larger whole. This includes component of, division of portion of, fragment of section of, and layer of" [21]. This definitions blurs the distinctions pointed out by [22]

Failure to make these distinctions is not merely of theoretical significance. For example, for component-integral object relations it is important to have a complete list of all relevant components (e.g. the contents of the thoracic cavity) so that properties of each component such as location and function can be specified; in contrast, for stuff-objects, it is only necessary to provide the user with information about each type of stuff', which for the heart is composed of muscle tissue or, at the microscopic level, muscle cells. The notion of a list of the individual cells that compose the heart is nonsensical–by definition, all cells of a certain sort have certain similar properties.

6.3 Neighborhood relations

A third type of relationship important in the anatomical domain are neighborhood relations, a semantic category in the UMLS hierarchy. The Digital Anatomist system uses two hierarchical relations: *branch-of* and *tributary relations*. However, relations which are not hierarchical, such as *near* and *underneath*, are needed for the navigation of anatomical images, especially 3D ones.

6.4 Associative relations

Because of the educational purpose of our system, associative relations such as 'function', 'site of' and 'definition of' are also important components of the knowledge base.

7 Conclusion and Future work

The availability of high quality, portable 3D models of anatomical entities offers the opportunity to create a 'virtual dissection' laboratory for medical students and faculty. In this lab, users can view subparts of the human anatomy separately, in relation to other parts, and in relation to a larger whole. The research we describe in this paper is based upon the hypothesis that a semantic network of anatomical information is a

AORTA CONCEPT SPACE (partial)

aorta :: is an :: artery

aorta :: has subparts :: [aortic arch, descending aorta,ascending aorta, aortic valve]

aorta :: originates in :: left ventricle

[vagus nerve, left recurrent laryngeal nerve] :: pass around :: aorta

[vagus nerve, recurrent laryngeal nerve] :: pass around :: aorta

ligamentum arteriosum :: connects :: [aortic arch, left pulmonary artery]

aortic arch :: has branches :: [left subclavian artery, left common carotid artery, brachiocephalic trunk]

[aorta, superior vena cava] :: pass through :: pericardium

pericardium :: envelops :: heart

heart :: has components :: [left atrium, left ventricle, right atrium, right ventricle]

subclavian artery :: passes over :: first rib

subclavian artery :: passes under :: clavicle

aorta :: has function :: [text explanation]

Figure 6

necessary preliminary to the development of a navigational system which reflects a conceputally sound organization of the material and is also computationally tractable. Based on the research discussed in this paper, we are planning to implement a prototype knowledge base with the following kinds of relations:

- An *is-a* hierarchy based primarily on that of the UMLS, as refined by the Digital Anatomist Program. However, the distinction between normal and abnormal entities will be shifted to lower down in the hierarchy so as to facilitate comparison of distinguishing properties of anatomical structures with particular abnormalities.

- At least three carefully defined *part-of* relations will be incorporated into the knowledge base: component-integral object, stuff-object, and place-area.

- A minimal number of neighborhood relations like *near* or *adjacent-to* will be included. The relations will not correspond to linguistic expressions, which are too ambiguous for our purposes, but rather will be formally defined.

An example of a partial concept space for the heart and aorta is presented in Fig.6.

In future work, we will discuss the formal characteristics of the relations we develop for the prototype, and we will report on the results of the formative evaluation of the usefulness of these relations for navigating anatomical information and of the applications that access the knowledge base.

References

[1] Ackerman M.J., "Fact Sheet: The Visible Human Project", National Library of Medicine, 1995.

[2] Spitzer V. et al.. The Visible Human Male: A Technical Report. *Journal of the American Medical Informatics Association (JAMIA)*, 1995.

[3] http://www.cs.stevens-tech.edu/vesalius

[4] Imielińska C., P. Molholt P., C. Bean, L. Laino-Pepper, E. Soliz, and R. Thumann, "The VESALIUS Project: Visualization of Lungs and Bronchial Tree Anatomy Lesson", *Proceedings of the First User Conference of the National Library of Medicine's Visible Human Project*, October 8-9th, 1996, Bethesda, MD 1996.

[5] Wacholder, Nina, Celina Imielinska, Stephen Johnson, Ewa Soliz and Pat Molholt, (1998) "Implications of 3D visualization for medical education", *Proceedings of the Medicine Meets Virtual Realit:6 (MMVR:6) Conference*, San Diego, California, January, 1998.

[6] Zucker, J., H. Chase, P. Molholt, E. Soliz and R.M. Kahn, "A comprehensive strategy for designing a web-based curriculum". In *Proceedings of the 1996 AMIA Annual Fall Symposium*, pp.41-46, Gabekt & Belfus, 1996.

[7] Kahn, R.M., P. Molholt, and J. Zucker. "CPM-Cnet: An integrated information and internet resource. In *Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care, SCAMC 1994*, pp.98-102, McGraw-Hill, 1994.

[8] Molholt P., C. Bean, C. Imielinska, and L. Laino-Pepper, "The VESALIUS Project: Creating a Network-Based Anatomy Curriculum", Technical Report No. 9605, Stevens Institute of technology, April 1996.

[9] Imielińska C., M. Downes, S. Hosakere, A. Khan, and W. Yuan,"2D Segmentation of Color Anatomical Data", Proceedings of the First User Conference of the National Library of Medicine's Visible Human Project, October 8-9th, 1996, Bethesda, MD; also accepted to a special issue of the *Journal of Computerized Medical Imaging and Graphics*.

[10] Kalvin, A.D., "Segmentation and Surface-Based Modeling of Objects in Three-Dimensional Biomedical Images", Ph.D. dissertation, NYU, March 1991.

[11] Lorensen, W.E., and Cline, H.E., "Marching Cubes: A High Resolution 3D Surface Construction Algorithm", *Computer Graphics: A Quarterly Report of SIGGRAPH-ACM* 21(4), pp.163-168, 1987.

[12] Brugman, Claudia (1988) *The Story of 'Over': Polysemy, Semantics, and the Structure of the Lexicon*. Garland Publishing, New York.

[13] Cimino, J.J., G. Hripcsak, S.B. Johnson, and P. Clayton. "Knowledge-based approaches to the maintenance of a large controlled medical terminology", *JAMIA*(1):35-50, 1994.

[14] Klavans, J., M. Chodorow, and N. Wacholder, "From dictionary to knowledge base via taxonomy". In *Electronic Text Research*, University of Waterloo, Canada, 1990.

[15] Schulz, Erich B., C.Price, and P.J.B. Brown, "Symbolic anatomic knowledge representation in the Read Codes Version 3: structure and application", *JAMIA*, 4:38-48, 1997.

[16] Lindberg, D.A.B., B.L. Humphreys, and A.T. McCray "The Unified Medical Language System" *Methods of Information in Medicine*, 32:281-291, 1993.

[17] Brinkley, J.F. and C. Rosse, "The Digital Anatomist distributed framework and its applications to knowledge-based medical imaging" *JAMIA*, 4(3):165-183, 1997.

[18] Rosse, C., M. Ben Said, K.R. Eno and J.F. Brinkley, "Enhancements of anatomical information in UMLS knowledge sources" in *Proceedings of AMIA*, pp. 873-877, 1995.

[19] Rosse, Cornelius, Jose L. Mejino, Bharath R. Modayur, Rex Jacobovits, Kevin Hinshaw and James Brinkley (1998) "Motivation and organizational principles for anatomical knowledge representation", JAMIA 5:17-40.

[20] Brachman, Ronald (1983) "What IS-A is and isn't: An analysis of taxonomic links in semantic networks" *Computer*, October 1983:30-36.

[21] UMLS Documentation (1997) 8th edition. http://www.nlm.nih.gov/research/umls/UMLSDOC.html.

[22] Winston, Morton E., Roger Chaffin and Douglass Herrman (1987) "A taxonomy of part-whole relations", *Cognitive Science* 11:417-444.