Chapter 2

Symbolic Modeling of Human Anatomy for Visualization and Simulation

A. Pommert, R. Schubert, M. Riemer, T. Schiemann, U. Tiede, K. H. Höhne

Institute of Mathematics and Computer Science in Medicine (IMDM)
University Hospital Eppendorf
Martinistr. 52, 20246 Hamburg, Germany
e-mail: pommert@uke.uni-hamburg.de

Abstract

Visualization of human anatomy in a 3D atlas requires both spatial and more abstract symbolic knowledge. Within our "intelligent volume" model which integrates these two levels, we developed and implemented a semantic network model for describing human anatomy. Concepts for structuring (abstraction levels, domains, views, generic and case-specific modeling, inheritance) are introduced. Model, tools for generation and exploration and applications in our 3D anatomical atlas are presented and discussed.

2.1 Introduction

Computer-aided education systems are often based on conventional *hypermedia* techniques, which offer a convenient access to logically related information. As a major drawback, however, these systems provide only a limited number of predefined texts or images, which often do not match the student's needs. User interaction remains basically the same as when studying a conventional book.

To overcome these limitations, we developed an interactive three-dimensional anatomical atlas which provides a "look and feel" close to a real dissection, based on a volume model [20, 23]. Volume visualization techniques are increasingly used in clinical applications, both for diagnostic and therapy planning purposes. While interpretation is generally left to the observer in clinical cases, applications in anatomy teaching require that the data structure also contains knowledge about *spatial* and more abstract *symbolic* properties of the objects represented in the image volume. In particular, a student must be able to access this knowledge in both directions:

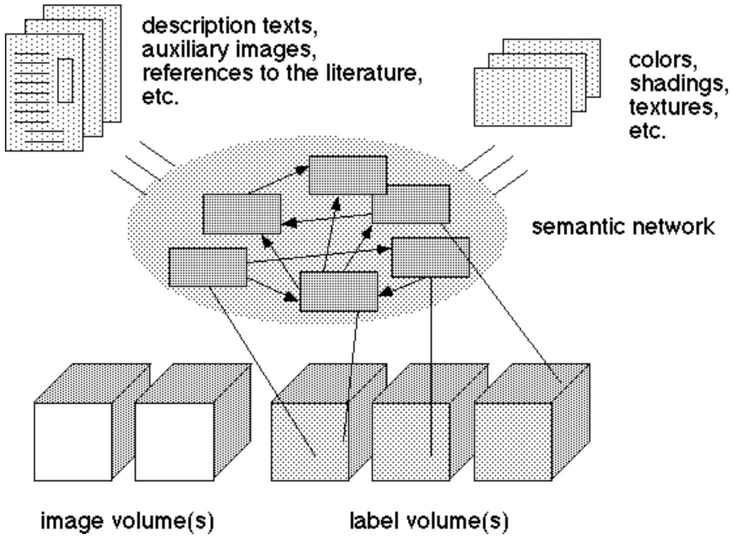

Figure 2.1: Basic idea of the "intelligent volume". Objects and links are described on a symbolic level using a semantic network. Their spatial properties are defined in some image and label volumes.

- access symbolic knowledge in the context of a 3D image (e.g. by automatically annotating, describing or painting a currently visible object)
- visualize symbolically defined 3D views (e.g. "show all gyri which are involved in a certain function")

For modeling on a symbolic level, we first implemented a rather simple model which provided only hierarchical part-of relations between different objects. While this model proved adequate for using the atlas as a visualization tool, it could not cover much of the information needed for student education. For example, relations between objects defined in different "views", such as morphology and function, could not be modeled.

In this paper, we are presenting the concepts introduced to extend the model towards capturing more of the complex structure of human anatomy. We also implemented these concepts into a knowledge base system, which is now being used as part of our atlas system. Experiences and future goals are discussed.

2.2 "Intelligent volumes"

As the basic link between spatial and symbolic knowledge, we developed a data structure which we call an "intelligent volume" [50]. Its two-level layout is shown in fig. 2.1. The lower level consists of one or more discrete data volumes, as obtained from computed tomography (CT), magnetic resonance imaging (MRI), or other imaging sources. A set of attributes is assigned to every voxel, indicating its membership to anatomical regions. These object labels are stored in one or more label volumes. The lower level is thus equivalent to the previously described *generalized voxel model* [19].

The upper level is a symbolic knowledge representation which contains further descriptive knowledge of the objects and the relations between them. In principle, any of the knowledge representation schemes developed in artificial intelligence could be used here [61]. However, since emphasis is on presenting the knowledge to a human observer, the applied formalism should be intuitive and easy to understand.

Therefore, we decided to use *semantic networks*, which have successfully been used for a wide range of applications. They also allow a comprehensive graphic representation.

2.3 Earlier work

In a previous paper, we reviewed a number of other projects for computer-based medical atlases [50]. None of them provided a symbolic model of the human anatomy suitable for 3D visualization. However, symbolic modeling of human anatomy is also a subject in some other disciplines such as artificial intelligence or medical documentation, which are thus to be considered.

In *computer vision*, propositional models have been developed e.g. for an automatic segmentation of MRI data [34]. The represented knowledge, however, consists mainly of topological descriptions of the anatomy such as "intracranial region is inside the skull". Clearly, these are of little use for our purposes.

Large models of the medical domain have also been developed in *expert systems* design [61]. Since the aim is to draw conclusions from a model, "if ... then ..." rules or similar descriptions are used. Again, this type of knowledge is not very suitable for our goals.

The methods developed for computer-based documentation of diseases, such as the systematized nomenclature of medicine [9] (SNOMED) or the the international classification of diseases (ICD), show some very different characteristics. These systems provide a very detailed set of terms (tens of thousands), together with individual codes. They are regularly updated to reflect the latest developments in medicine. In SNOMED, a disease is coded using seven individual dimensions describing e.g. its cause and location. However, as mere collections of terms, these systems do not provide descriptions of the relations between various terms, as needed for modeling the anatomy. Furthermore, everything is centered on diseases, so that other aspects are quite poorly covered.

2.4 Generic modeling

In our semantic network approach, objects and their properties are modeled as nodes and their relations as links. Some attributes of an object are:

- names (preferred terms, synonyms, colloquial terms) in various languages
- pointers to related medical information (description texts, images, references etc.)
- visualization parameters such as color, shading, and texture.

Objects may be anatomical objects, such as the *left precentral gyrus*, or abstract concepts such as *light stimulus* or *vision*. For choosing preferred terms and synonyms, we built on standardized anatomical nomenclature as far as possible [26]. Since arbitrary numbers of terms can be defined, it is thus easily possible to find a preferred term for a colloquial, or to get a list of translations.

Some of the currently defined link types along with some properties[1] are listed in tab. 2.1. They include some relations which can be found in many semantic network models, such as part hierarchy (*PartOf*) or generalization (*IsA*), as well as some more specialized links, such as stimulus propagation (*PropagatingTo*). The resulting topology is a real net, since any node may have numerous parents and children; even circles appear e.g. when modeling the vascular system. So far, we are quite conservative in introducing new link types, partly because the semantics of the net have to be carefully defined if several link types are involved.

[1] Be $R \subseteq A \times A$ a relation. A is irreflexive: $\neg(aRa)$; symmetric: $aRb \Rightarrow bRa$; anti-symmetric: $aRb \Rightarrow \neg(bRa)$; transitive: $aRb \wedge bRc \Rightarrow aRc$; $\forall a, b, c \in A$.

			symmetric	anti-symmetric	transitive
object	*PartOf*	object	—	×	×
object	*IsA*	object	—	×	×
object	*PassingThrough*	object	—	—	—
object	*PropagatingTo*	object	—	—	×
function	*DependingOn*	object	—	×	—
artery	*BranchingFrom*	artery	—	—	—
artery	*AnastomosingWith*	artery	×	—	×
object	*SuppliedBy*	artery	—	—	—

Table 2.1: Some link types and their properties. All link types are irreflexive.

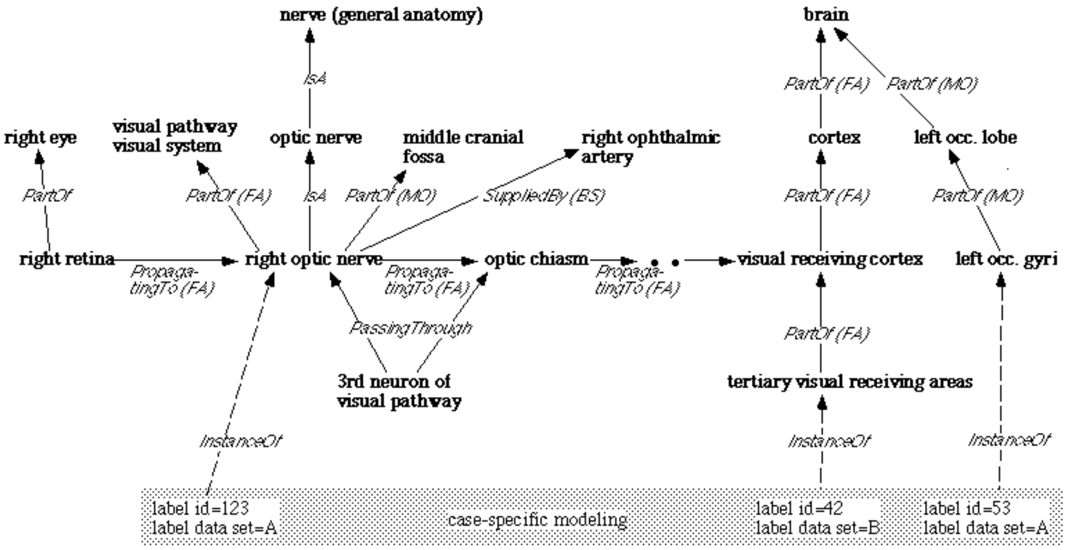

Figure 2.2: Detail of a generic model, showing the right optic nerve and some of its relations. Some links are valid for certain domains only, such as morphology (MO), functional anatomy (FA), or blood supply (BS). The case-specific model provides the link to a spatial description.

All properties which can be modeled with the concepts introduced so far are (usually) valid for the human anatomy in general. The same knowledge may thus also be used for other atlases, based on different data sets. We call this part of the semantic network the *generic model.* The described knowledge is thus actually about *classes* of objects and links.

2.4.1 Structuring medical terminology

For a tool like an atlas which is to be used in student education, the vast medical domain must be presented in a well-structured way and in comprehensible quantities. With no further elements for structuring the knowledge available, the net soon becomes very confusing. For example, the *precentral gyrus* and the *motor cortex* are both *PartOf* the *cortex.* However, since both terms belong to different "views", it is obvious that these should also be reflected in the model. The rest of this section is thus devoted to techniques for structuring the medical terminology. An accompanying example is shown in fig. 2.2.

Abstraction

A natural approach to structuring the knowledge is to look at the various subjects taught in medical education. One of these concepts is to divide anatomy into *general anatomy* (example: *gland*) which describes general principles, and *special anatomy* (example: *parotid salivary gland*) which describes individual aspects. Terms on different levels of abstraction are connected using *IsA* links. The degree of abstraction can be modeled as an attribute of an object. Example:

$$\begin{array}{rcl} spongiosa_{general} & PartOf & bone_{general} \\ long\ bone_{general} & IsA & bone_{general} \\ epiphysis_{general} & PartOf & long\ bone_{general} \\ femur_{special} & IsA & long\ bone_{general} \\ femur_{special} & PartOf & skeleton_{special} \\ right\ femur_{special} & IsA & femur_{special} \end{array}$$

From this follows e.g. that the *right femur* has an *epiphysis* (it is a *long bone*) and a *spongiosa* (as every *bone*).

Domains

Another important classification reflects division of anatomy into different *domains* such as morphology, functional anatomy, and pathological anatomy. However, just like some lectures focus on a specific subject, it is often useful to define even smaller domains such as *arterial system*, if this is to be modeled in greater detail.

Abstraction levels and domains are *orthogonal* (independent) dimensions. An anatomical object may be considered from different points of view, such as

$$\begin{array}{rcl} optic\ chiasm & PartOf_{morphology} & brain \\ optic\ chiasm & PartOf_{functional\ anatomy} & visual\ system \\ cornea & PartOf_{morphology, functional\ anatomy} & eye \end{array}$$

Clearly, domains have to be modeled as an attribute of a link (fig. 2.2). The latter example shows that a link may be valid in more than one domain.

It should be noted that limiting a link to be valid in certain domains only does not mean to introduce a new link type; the basic semantics of the link types such as *PartOf* remain unchanged. This point is very important for implementation. New domains may thus freely be defined at atlas run-time.

Other orthogonal dimensions

In the same way as abstraction levels and domains, additional orthogonal dimensions can be introduced. This may be required to make things more obvious, e.g.

$$\begin{array}{rcl} hypophysis & PartOf_{morphology-topographic} & sella\ turcica \\ hypophysis & PartOf_{morphology-systematic} & diencephalon \end{array}$$

Other orthogonal dimensions are important for modeling natural *variation*, such as sex or age (fetus, child, adult etc.). A further dimension may be introduced to distinguish between different sizes of the considered structures (macroscopic, microscopic, molecular etc.) [3]. The latter may be important e.g. to keep the descriptions of functional systems consistent, which are modeled on various levels of resolution.

Views

Selecting a set of abstraction levels, domains, other orthogonal dimensions and link types, a student working with the atlas can choose a subset of the knowledge base which is relevant for a particular task. We call this a *view*. All other parts of the network are not considered and will thus not be visible. A typical view for a beginner may be *special anatomy/morphology and functional anatomy/macroscopic anatomy*.

2.5 Case-specific modeling

Re-considering our "intelligent volume" approach, we now have a generic model of object and link classes on a symbolic level. The object classes are not yet linked to the object *instances*, whose spatial properties are described in some image and label volumes. The link is provided by a symbolic description of an instance, which has the following attributes:

- reference to a generic object (*InstanceOf*, see fig. 2.2)
- label ID, intensity range, data set names and interpretations (e.g. intensity, label, radiation dose)
- case ID
- optionally all attributes which can be defined for a generic object.

The case ID is used to distinguish between different sets of image and label volumes, e.g. representing an atlas reference case and one or more patient cases. They may thus be viewed and manipulated simultaneously [47]. The optional parameters may be used to overwrite the corresponding attributes defined for an object class (see next section). As opposed to the generic model, a case-specific model thus consists of a number of data volumes and a symbolic description of the instances.

2.6 Inheritance

Many attribute values of an aggregated object (e.g. the display color of the *brain*) should be the same for all its parts (e.g. the *gyri*), unless something different has been specified (e.g. a specific color for *white cerebral matter* and its parts). Instead of assigning the same value to every object, we are using a default inheritance mechanism [61] along the *PartOf* and *IsA* links. Object description is thus considerably simplified.

Due to the network topology, however, two or more different values may be inherited over different paths. In this case, the value is taken from the node which can be reached using the smallest number of links. The results of this algorithm meet our intuitive expectations quite well. In the (very rare) case where two different values are inherited over two paths of the same length, a value can be assigned directly to the object.

As a further extension of our inheritance algorithm, attribute values defined in the generic model may be modified for individual cases, using a *Z inheritance* scheme [61]: generic object and corresponding instance are considered as pairs, with the value taken from the instance (if defined), otherwise from the generic object (if defined), or otherwise inherited from a superior generic object/instance pair, following the above described principles.

2.7 System environment

Our 3D atlas environment consists of the four main modules 3D visualizer, knowledge base system (KBS), user interface, and interactive segmentation. The first three are making the actual atlas program VOXEL-MAN/atlas [23]. While the user interface is dealing with objects in terms of symbolic descriptions, the visualizer knows object instances only in terms of label IDs and data sets. The KBS which stores all generic and case-specific knowledge on the symbolic level provides the link between these two levels. Symbolic definitions may be saved to and read from file, using a simple language which can easily be manipulated.

The program INTERSEG [46] used for interactive segmentation is also fully integrated and has access to visualizer and KBS. After an object has been defined, its spatial extent is saved in a label volume with an automatically chosen label ID, which is assigned to every voxel contained in the object. If the object is already described in the generic knowledge base, it is sufficient to select its name from a list in order to link all available symbolic knowledge to the new spatial definition. The symbolic description is written to a case-specific list of instances. If the object is still missing in the generic knowledge base, its symbolic description can be added now.

For implementation of the atlas system, we put emphasis on portability and independence of other software products. Therefore, we used only standard tools such as ANSI-C and OSF/Motif. VOXEL-MAN/atlas is currently running on various workstations (DEC, SUN, HP, SGI). As regards the KBS, we abandoned using readily available data base or knowledge representation systems also for reasons of speed and flexibility.

2.8 Applications

Using the above described methods, a generic description of the human anatomy including all parts of the atlases previously created at our institute (e.g. brain, skull, abdomen, pelvis, fetus) has been defined. This still very incomplete model serves as a frame which is continuously updated. Current projects include modeling of bones and muscles of the whole human body, a very detailed modeling of the visual system and the vascular system of the head, and modeling of a "microscopic domain" (different tissue types etc.). New generic objects are usually defined after a corresponding instance has been segmented in a data set; however, a broad coverage of e.g. a functional system requires that many objects have to be modeled which are (far) below the resolution of current tomographic image data. As yet (June 1994), the generic model contains about 1200 nodes, 2400 links and 3300 names.

Several views of the generic knowledge base are shown in fig. 2.3. These representations may be called during interactive segmentation or during an atlas session. Every view provides only a part of the whole symbolic model, according to the choice of abstraction levels, domains, and link types. Objects may be selected by simply clicking to them. For searching, the KBS provides an additional fault-tolerant interface which also allows wildcards.

2.8.1 Visualization

In the following, it is shown how the symbolic knowledge modeled in an "intelligent volume" may be accessed in a pictorial context during an atlas session. Emphasis is on showing aspects which are related to semantic network modeling; for an extended presentation of the full atlas user interface and applications of VOXEL-MAN/atlas, we refer to [51, 56].

As a consequence of the space-filling model, the contents of the knowledge base may be accessed at any point of a 3D image. Symbolic descriptions may thus be obtained by simply clicking on an image. Object

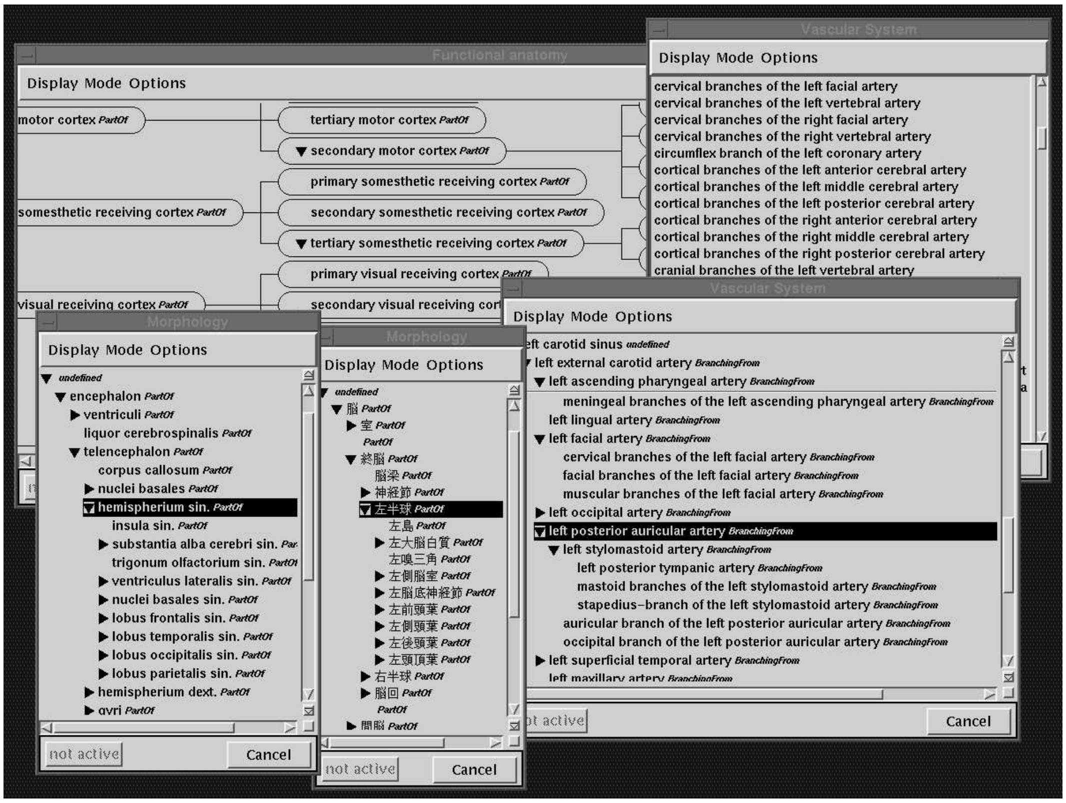

Figure 2.3: Several views of the knowledge base, showing selected parts of the whole generic model. Form of presentation (alphabetical list, structured list, subtree of the network) and language may be chosen.

names appear on a pop-up menu, together with the names of the objects they are related to, ordered by domains and link types (fig. 2.4). By selecting a name, objects may automatically be annotated, painted, added or removed, or additional texts and other information may be requested. In the shown example, the chosen object (*right optic nerve*) appears several times, according to the different domains it is defined in (fig. 2.2). To ease comprehension, individual background colors are used for all domains.

Even though the optic nerve in fig. 2.4 may be seen in different domains, according to the symbolic modeling, its spatial extent is defined using only one label ID / label data set pair. Several such pairs are needed if the objects defined in different domains have different, but overlapping spatial descriptions, such as morphological and functional regions and blood supply areas of the brain. On the pop-up appear thus several different names in the different domains (fig. 2.5, see also fig. 2.2).

Of course, all functions of the atlas may also be called from the knowledge base interface shown in fig. 2.3, even for objects which are not currently visible.

2.8.2 Simulation

The model can also be used to simulate and visualize functional systems. For example, a part of the visual system is described as follows:

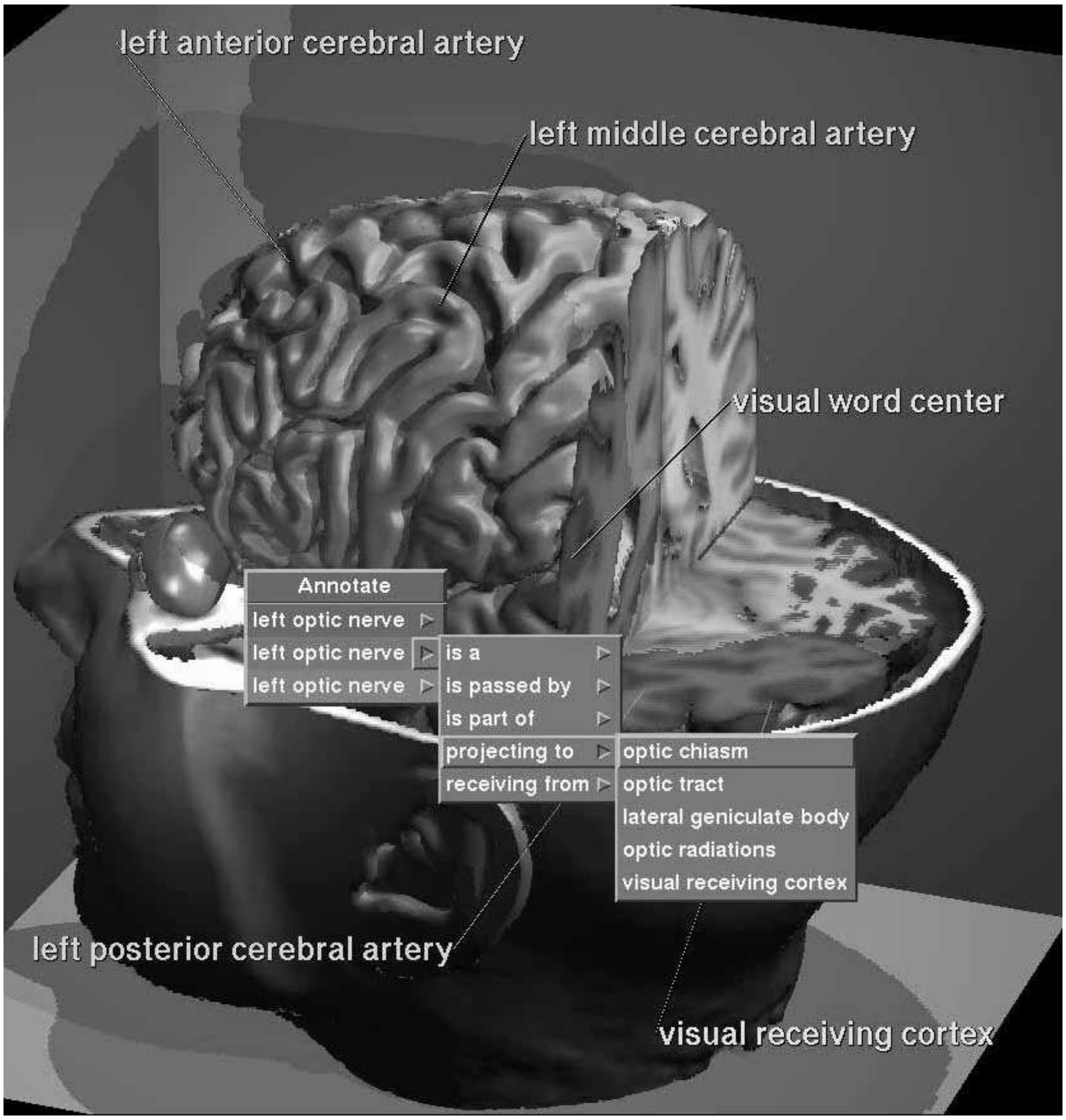

Figure 2.4: Pop-up menu, invoked by clicking on an atlas "camera". The *right optic nerve* appears three times, indicating the domains it is defined in (fig. 2.2). The student may select the domain and the link type he is further interested in.

right retina	*SuppliedBy*	*right central artery of the retina*
right lens	*PropagatingTo*	*right retina*
right retina	*PropagatingTo*	*right optic nerve*

Stimulus propagation is shown on a 3D image following the *PropagatingTo* links by subsequently painting the involved objects. Furthermore, if an object in this path is "destroyed", e.g. by breaking its blood supply, the path thus breaks up into an active and an inactive part, which can be shown in different colors.

It is a decisive advantage of our approach that the knowledge described in the model can be visualized on a 3D image. For the future, it might also be interesting to extend the atlas to include a (simple) expert

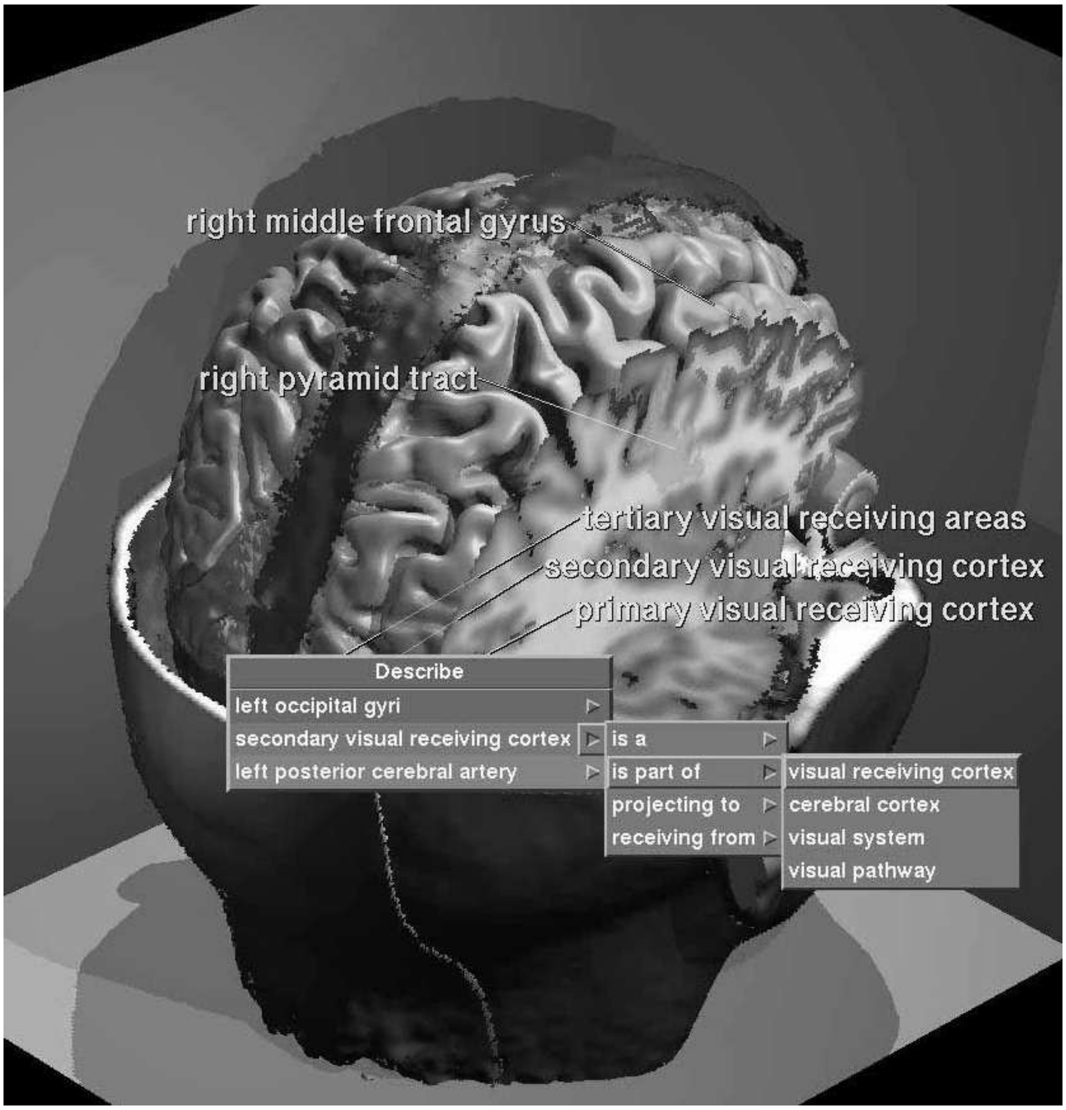

Figure 2.5: Another pop-up menu. In this case, the investigated point in 3D space represents three different objects with different spatial extents, defined in different domains (fig. 2.2).

system [41]. For example, the states of the objects along a *PropagatingTo* path and possible consequences hereof could be evaluated. The other way round, the possible causes for the loss of a function could be investigated and visualized. It should be noted, however, that this is not only a matter of symbolic modeling. As has been shown, some properties are defined on the spatial level only, e.g. whether a functional region and a blood supply area overlap. In order to automatically evaluate the consequences of a break in blood supply, some "spatial reasoning" capabilities will thus be required.

2.9 Conclusions

In this paper, we presented an approach for modeling the human anatomy on a symbolic level. In contrast to other models which aim to provide knowledge for machine-based interpretation, our model is to be used mainly for presenting knowledge about the anatomy to a human observer. We therefore introduced several concepts for structuring the knowledge such that it may be accessed in comprehensible quantities. It was shown how knowledge on a symbolic level is linked to spatial knowledge and utilized in an interactive 3D anatomical atlas.

We still consider ourselves to be in an early stage of this project. Due to the flexibility of the "intelligent volume" model, there is a broad range of possible future developments. Considering the knowledge base, we currently see three major directions. First, the model has to be filled with many more nodes and links. This is not only a matter of quantity, but also of growing complexity, which will show where the modeling has to be improved. As we pointed out, there are several projects currently going on. Since different people are involved, keeping the whole model consistent is a major concern.

Second, better modeling tools are needed. As the generic model gets more complex, it should be possible to edit the net structure with a graphic editor. Likewise, automatic consistency checking is not yet sufficient.

Third, it is important to more formally define the semantics of the model, especially if several link types or even "spatial reasoning" are involved. This is also a prerequisite for using the model for advanced functional simulations.

Acknowledgements

Examples are in cooperation with G. Weske and R. Cosack (both IMDM), who are currently modeling visual and vascular systems, respectively. We are grateful to S. Eiho (Univ. of Kyoto), who very patiently translated and coded a large number of Japanese terms. Additional advice was kindly provided by M. Shimizu (Univ. of Kyushu at Fukuoka). The MRI raw data set of the head is courtesy of Siemens Medical Systems, Erlangen.

Suggested further reading

Symbolic modeling is also subject of several other atlas projects. In the Digital Anatomist, the Unified Medical Language System (UMLS) is extended for modeling gross anatomy and neuroanatomy [44], with a number of applications illustrated in [7]. In the SAMMIE project, a frame-based approach is taken [35].

As a remaining problem, the somewhat unclear semantics of partonomies (*PartOf* relationships), especially regarding transitivity, has been identified [2, 52]. Some implications and solutions developed in the VOXEL-MAN project will be presented in [49].