Chapter 1

A New Representation of Knowledge Concerning Human Anatomy and Function

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

Institute of Mathematics and Computer Science in Medicine
University of Hamburg, Germany

Abstract

By integrating concepts of computer graphics and artificial intelligence novel ways of representing medical knowledge become possible. They allow unprecedented new possibilities ranging from 3D interactive atlases to systems for surgery rehearsal.

1.1 The abundance of medical information

No other subject has probably been investigated in more detail than the human body. A vast amount of data has been accumulated over centuries in the qualitative form of text and pictures. With the advent of the exact sciences, to a limited extent, formal analytical descriptions e.g. in physiology were included. The speed at which new data are created is still increasing. The information has been accumulated in myriads of books and journal articles organized in a large number of subfields some of which often do not take account of each other. Quite recently so-called multimedia systems have been built which allow the fast and random access to text and pictures. Yet the underlying paradigm of such systems is still that of the old book with pages consisting of a static combination of textual and pictorial symbols and, to a small extent, the equations that describe various processes and functions. The entire organisation follows the rules imposed by the technique of printing. Each book or article describes a specialist's view of a specific topic for a certain audience. One and the same body part would be described in different ways, both textually and pictorially, depending on the exspected interest of the readers - structural or functional, diagnostic or surgical, for students or experts.

However, the state of the art of computer science suggests that knowledge about the human body could be organized in a more condensed and abstract way. Instead of dealing with plain text, knowledge engeneering allows the structuring of symbolic information in a much more consistent and flexible way. Instead of dealing with static images or movies, computer graphics allows for the creation of realistic interactive 3D models of human anatomy and function. Instead or in addition to a textual or formal description, dynamical processes can be simulated and visualized. If these concepts could be integrated into a single more general description of the human body, this would enable us to potentially extract all classical knowledge representations for different purposes as views to this abstraction. More important, we could take the word 'view' literally in that we could access the information via a visual representation of the model.

It is clear that the generation of such an abstraction is a huge undertaking. Perhaps it will not be feasible in full completeness at all. However, it seems to be possible to condense its spatial and symbolic description in a hybrid data structure, which has the character of a model. As shown with an interactive 3D atlas of the human gross anatomy [20, 56, 51], based on such ideas this concept is feasible.

1.2 Early attempts to formalize anatomical knowledge

First efforts in developing digital representations of anatomical knowledge resulted in data bases, which enabled retrieval of relevant items for a limited part of anatomy. These systems were mostly local solutions, that enabled researchers to archive both their own results and related results described in the literature. An overview of concepts used for constructing such databases can be found in [3]. Although these databases were an improvement for research, they were deficient in connecting descriptions to anatomical structures and visualizing the stored information in a "natural" way.

The next step was to include graphics and images which can be found in standard books and atlases like the *Talairach-* or *Schaltenbrand-Atlas* [55, 45]. Usually hyper- and multimedia techniques are used to link symbolic with pictorial data. Some of these systems include histological or photographical images from cadaver dissection. Representations of this concept are *NeuroDatabase* [60] or the data base of the BrainMap project[14]. On the same basis teaching programs have also been developed such as *A.D.A.M* [31], or *BrainStorm* [11]. These systems enable a convenient access to different anatomical information, but the stored images are precomputed views and, therefore, a "look between the atlas pages" is not possible.

Digital anatomy representations have also been described which were designed to facilitate interpretation of images used for clinical diagnosis [13, 16, 28, 29, 38]. They have in common that (a) anatomical structures are represented as sets of coordinates or 2D-contours only, and (b) the anatomical knowledge is restricted to a (sometimes hierarchical) list of anatomical names for a limited number of morphological objects. Likewise purely abstract descriptions whitout a pictorial representation have been described for automatic image interpretation with methods of computer vision [8, 43]. Only a few digital anatomy representations based on a 3D model have been described [33, 6, 59, 15]. These projects have in common that objects are represented as 3D contours only and are therefore hollow, which does not correspond to reality.

A first step to acquire detailed data for a space-filling model is proposed in the "Visible Human" project [1, 54], where data of two cadavers are being acquired, respectively, by Magnetic Resonance (Tomography) and Computed Tomography, and by photographs of anatomical cross-sections. As yet, the project does not define a procedure on how to convert the data into an anatomy model.

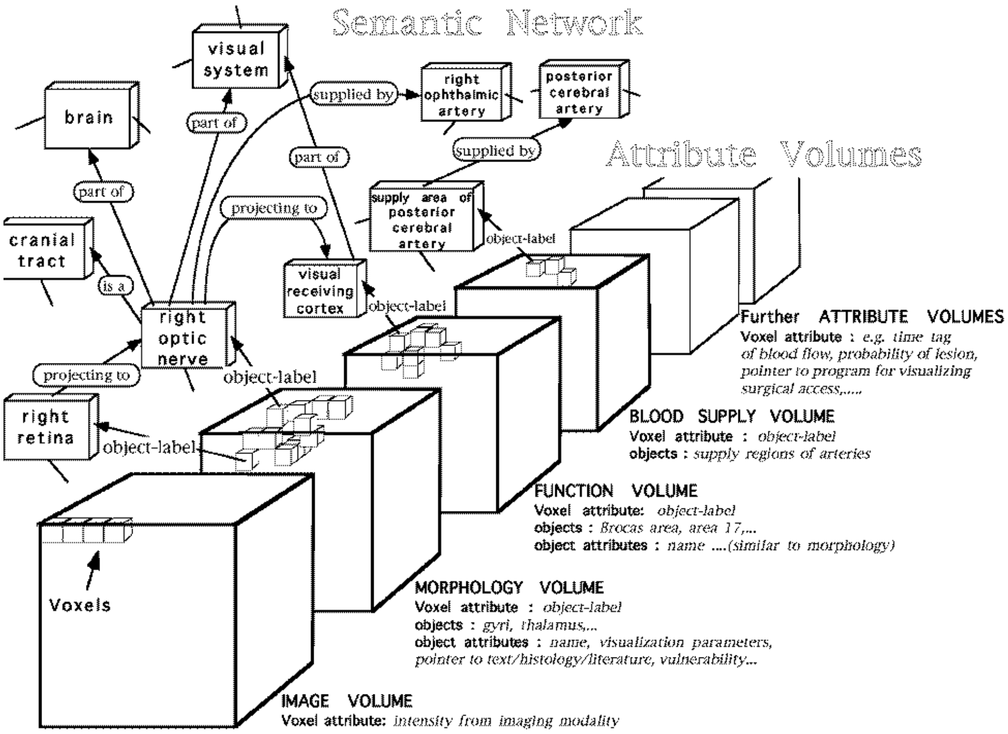

Figure 1.1: Block diagram of the knowledge representation consisting of a semantic network linked to a digital volume representation.

1.3 "Intelligent volume" approach

1.3.1 Basic idea

The key idea underlying the new approach is to combine in one single framework a detailed spatial model enabling realistic visualisation with a symbolic model of the human body (Fig. 1.1). The spatial model is sampled from a living subject via cross-sectional imaging modalities such as Computed Tomography, Magnetic Resonance Tomography or histological sections.

This volume is segmented into its constituents in as much detail as possible. These constituents will differ for the different domains of knowledge like structural and functional anatomy. One and the same voxel may belong to different voxel sets with respect to the respective domain. The membership is characterized by object labels which are stored in "attribute volumes" congruent to the image volume (Fig. 1.1). Further attribute volumes may be added which contain e.g. the incidence of a tumor type or a time tag for blood propagation on a per voxel basis.

The objects themselves bear attributes as well. They may be divided in two groups: attributes defining their visual appearance, like color, texture, reflectivity; and attributes concerning their meaning like names (in different languages), pointers to text or pictorial explanations, but also features like vulnerability or mechanical properties, which might be important e.g. for surgical simulation.

So far the model describes a spatial distribution of objects but not their interrelations. For a relational description among different possibilities offered by artifial intelligence [61] the technique of semantic networks was chosen [43, 37]. This allows the expression of links between objects via relations with given semantics. Three groups have been considered so far: those for modelling abstractions; for structure; and for functions. The following examples outline the essentials of the semantic network model (refer to [42] for details).

1.3.2 Modeling descriptive information with semantic networks

Fig. 1.1 shows a small section of the network containing some relations concerning the right optical nerve. Its structural relations are modelled via the Part Of relation. Since it is possible to view structures from different points of view (which are typically the different knowledge domains), they have the knowledge domain as attribute. Thus concerning morphology the right optical nerve is a part of the brain, while in functional anatomy it is part of the visual system.

Functional relations such as stimulus propagation may be modelled via the ProjectingTo relation: a stimulus is propagated from the retina through the optic nerve finally to the visual receiving cortex. The fact that the right optic nerve is supplied by the right ophthalmic artery is described with the SuppliedBy relation. Finally, the abstraction saying that it belongs to the class of cranial tracts is expressed by the IsA relation.

From a model containing these relations, we could ask about the position of anatomical objects in a structural hierarchy, we could ask what structures/areas are involved in stimulus propagation and, beyond that we could enquire about consequences of a failure of the right ophtalmic artery.

Despite of the elegance of the approach it must be clear that today's medical knowledge is not sufficiently consistent to allow a universally agreed description. Thus only part of the available knowledge can be represented in the highly structured form of the semantic network linked to a voxel representation. The rest has still to be in the form of text and classical pictures which is, however, linked to the respective objects in the semantic network via pointers (see Fig. 1.1). This hybrid structure of the model does not exclude unstructured information, but allows continuous conversion of poorly structured data like plain text into the higher structures of the model. The generation of the model is therefore a continous process driven by practical experience with its application.

1.3.3 Anatomy model of the VOXEL-MAN system

With the tools above described a semantic network model of the body has been built, which contains so far 1000 objects with 2500 links in in the domains morphology, pathology, functional anatomy and vascular system. It is being refined continously. It is called the generic knowledge base, since it contains general knowledge which is independent of an actual spatial visual representation (e.g. a head acquired with Magnetic Resonance Tomography). For linkage with a visual representation via an actual specimen the applicable items of the knowledge base are linked to the respective image volume. The resulting data structure is called intelligent volume. Practical experience has shown that a user easily looses track if the full semantic network is available at any time. Thus the accessible knowledge is restricted to what is necessary for an application to a subset, which is called a view.

The actual generation of the intelligent volume is shown with the example of the atlas of the human brain and skull. To provide a spatial representation of the brain an image volume of an individual was taken from Magnetic Resonance Tomography with resolution of 1.5 mm, the skull was sampled from Computed Tomography with an isotropic resolution of 1 mm. The segmentation was done in the following way: Initially the gross constituents, such as skull, brain, and ventricular system, are segmented by an interactive segmentation program. An expert selects the appropriate regions using the intensity window control and special image processing tools [21, 46]. The subdivision of the coarse components into detailed structural regions (like the gyri) is done by an expert using a "volume editor". For example, a gyrus would first be identified by painting its extent on the 3-D brain surface image. Gray matter (voxels below the painted surface) is identified by specifying an intensity threshold. Refinements can now be made by repeating the painting procedure from different viewing angles and finally on the cross-sectional images. Objects not present in the image data such as nerves are added via a "graphics editor". Once the spatial extent of an object is defined, the link to the generic knowledge is established by giving it a name. If, for example, an object has been assigned the name "precentral gyrus", attributes and relations are attached to the object automatically. As one of the consequences, it receives the standard visualisation parameters (color, reflectivity, ...) of "cortex". Even if advanced editing tools are used, this procedure of generating an intelligent volume is tedious. For the atlas of the skull and brain it took in the order of one man year to segment the 300 anatomic objects available so far.

1.3.4 The new features of the approach

However, once the work is done, a knowledge representation is established, which allows the user to freely navigate in both the spatial and the descriptive world. For visualizing the spatial world, computer graphics provides all effects that a painter would use to paint a 3-D scene onto a planar screen just by applying simple laws of projective geometry and optics [19, 57]. In the resulting system a viewer can generate arbitrary perspective views by choosing a certain view point, focal length and light direction of virtual cameras.Their pictures appear in windows on the screen.

For exploration of the model a user may address an object he is interested in either by clicking on its spatial representation in the perspective images or on a symbol in a textual or graphical representation of the knowledge base. The system allows the user pictorially to add/remove objects, color-mark related objects, and to show corresponding histological sections. It also provides information as subnets describing relations of the object (as seen in Fig. 1.2) or a textual description. The VOXEL-MAN system is presently beeing validated in 25 institutions around the world. For more details of the implementation we refer to [56].

1.4 Novel possibilities evolving from "Intelligent Volumes"

The result of the approach is a model represented by a data structure, which is primarily independent of an application. We can, however, derive very different applications from it just by providing appropriate exploration tools. Thus the new kind of model can be used in the following rather different ways.

1.4.1 Interactive exploration for learning anatomy and function

In the case of a 3D atlas of the brain [25] a student could begin to explore the brain model with an outer surface image of the head. As the first step, he or she could simply make a cut across the head and explore the resultant section. However, an anatomist might prefer to remove the tissue layer by layer. This can easily be simulated by selectively removing objects in the region delineated by the cut planes as shown in Fig. 1.2. Since the cut surfaces exhibit the texture of the interior, the scene has the appearance of an actual dissection.

The crucial advantage of the computer model is that, for any visible voxel, all related information can be queried. As seen in the figures, the corresponding object names in different domains (such as the name of a gyrus or a cortical region) can be annotated at the cursor position, whether there are sections or surfaces. The corresponding region can be colored as well. The language for the annotations can be

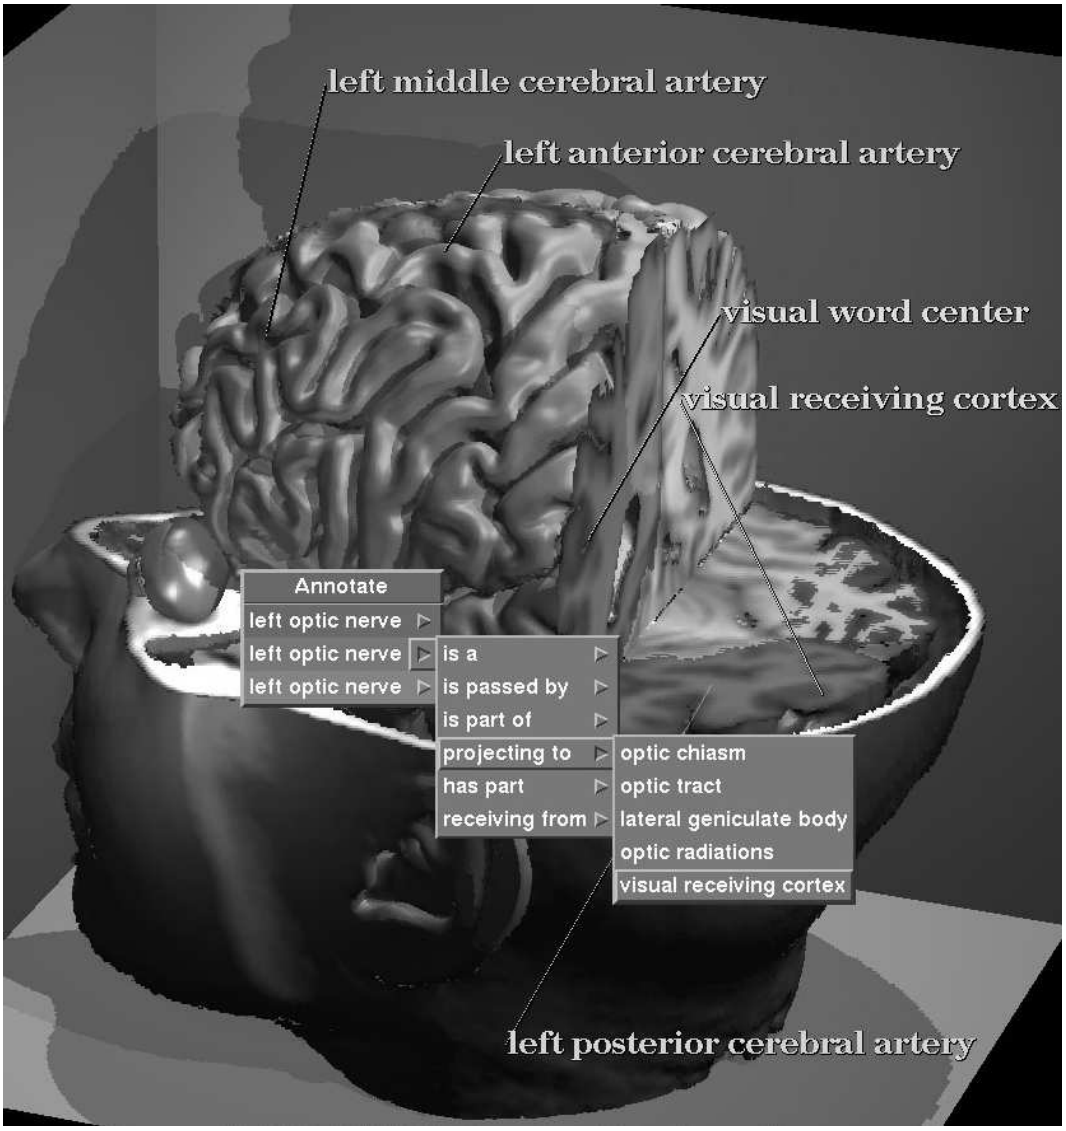

Figure 1.2: Exploration of the brain. Arbitrary cutting reveils the interior. The user has accessed the information available for the optic nerve concerning functional anatomy, which appears as a cascade of popup menues. He has asked the system to color-mark some blood supply areas and cortical areas. He can derive from the colors that the visual word center is supplied by the left middle cerebral artery.

chosen among English, French, German, Japanese and Latin (see also Fig. 1.3). In the case of brain atlas text descriptions and sample histological sections corresponding to the objects may be accessed.

Questions about the relationships between basic regions can be answered as they are described in the knowledge base. In the pilot project, for example, the structural "parents" of a basic region (e.g. the lobe to which a gyrus belongs) or functional regions traversed by a stimulus can be queried and colored, displayed, removed, etc. (Fig. 1.2). Relations between objects in different domains may in many cases be explored in a simple pictorial way. If, for example, we want to know which cortical regions would

Figure 1.3: When the model is derived from a radiological cross-sectional technique the radiological cross-sections (here from Magnetic Resonance Tomography) may be viewed in the context of basic anatomy. Any voxel may be accessed, even in different languages (bottom left).

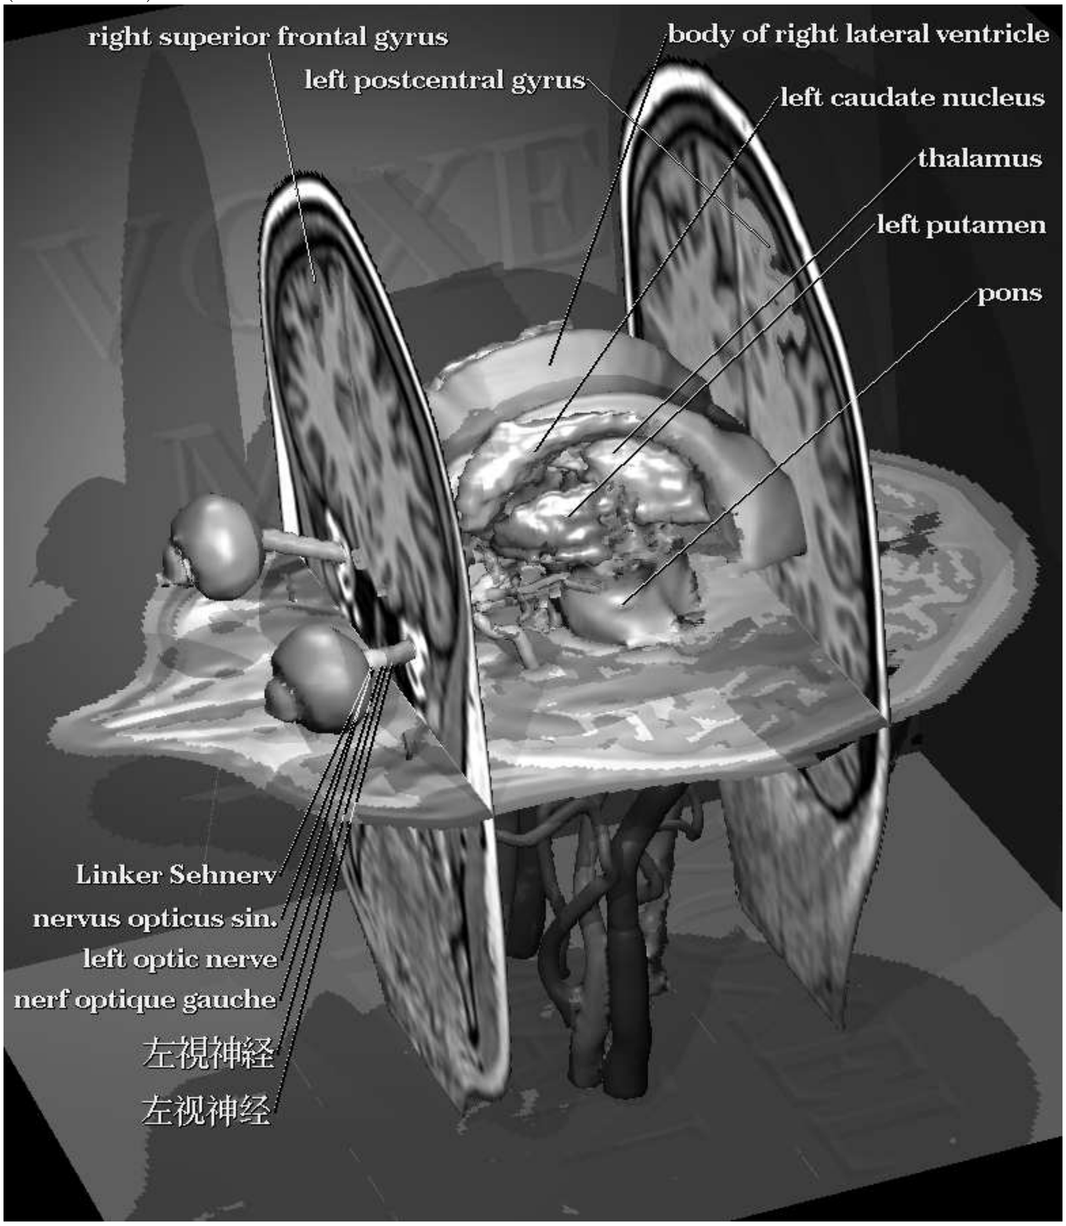

be affected by a stenosis in the middle cerebral artery, we would just ask the system to mark its supply region on the actual image an then point to the included regions to obtain the color marked functional regions (Fig. 1.2).

We can in the reverse case browse through a graphical or textual display of the knowledge base and select an object or group of them that can be treated as shown above. For example, we could ask for the exclusive display of the thalamus and neighboring objects as shown in Fig. 1.4. In this case, it is

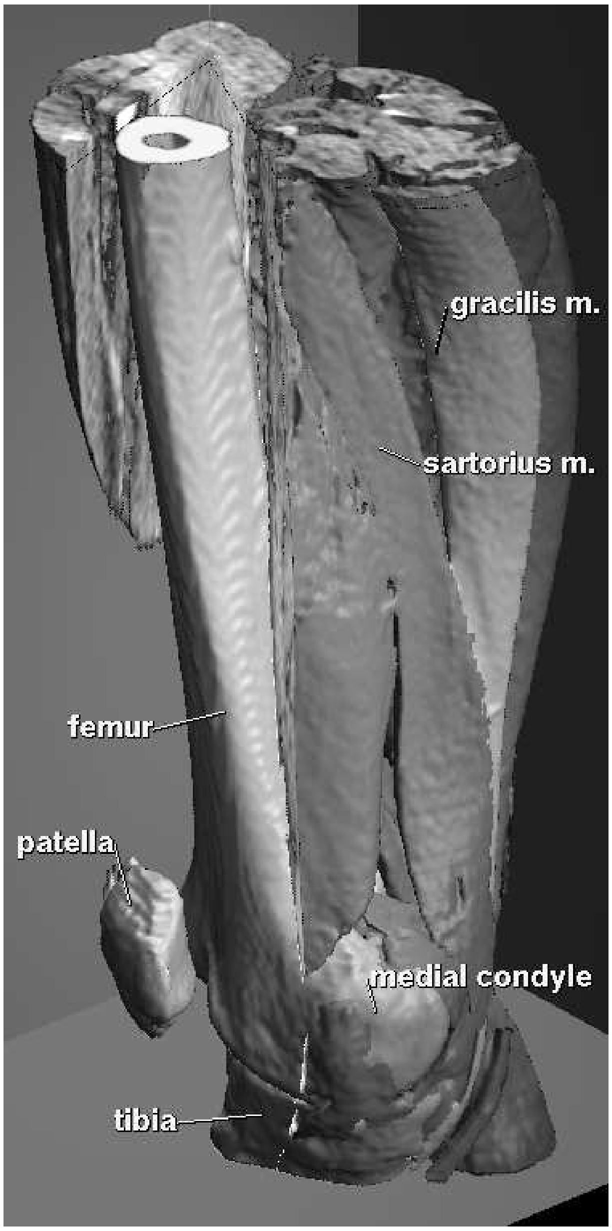

Figure 1.4: View generated with the 3D interactive atlas of the musculo-skeletal system derived from Computer Tomography.

advisable to add the ventricular system for orientation, which, unlike in an actual dissection, can be displayed as a solid object. While the brain model is the most elaborate one so far, other atlases are being generated such as one of a fetus (Fig. 1.5) and the musculo-skeletal system (Fig. 1.4).

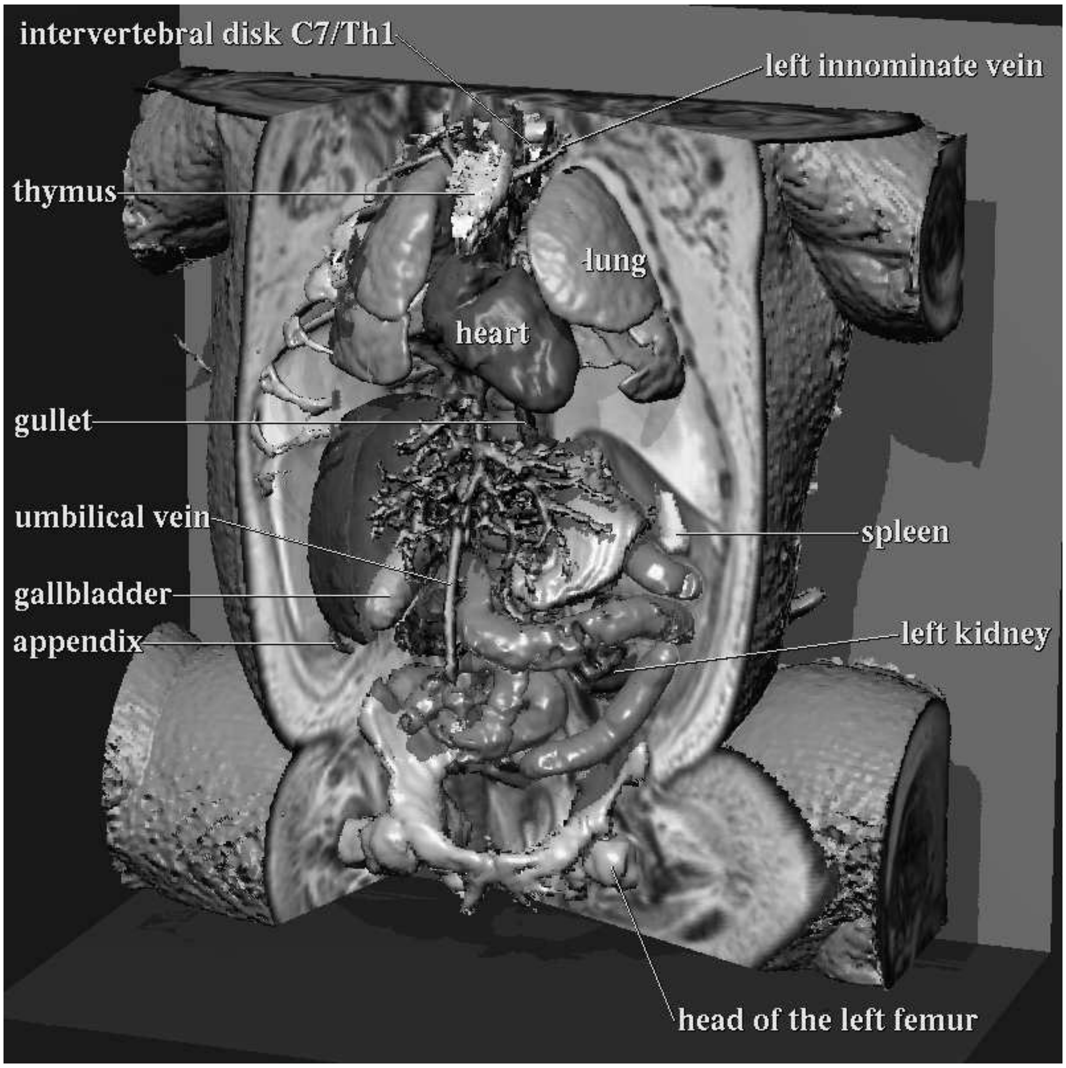

Figure 1.5: View generated with the interactive 3D atlas of a human fetus, derived from Magnetic Resonance Imaging.

1.4.2 Interactive reference system for radiological image interpretation

Since derived from radiological imagery, the model lends itself for teaching and reference purposes in radiology. As shown in Figure 1.6, this can be done by presenting cross-sectional images in the context of basic anatomy. Even more, when the visual model is derived from x-ray computed tomography, it is easy to simulate an x-ray projection from a chosen direction with any beam geometry. Since we know the objects that cause the projection image, we can determine the contribution of the individual objects to the final intensity at any location (Fig. 1.6).

1.4.3 Interactive surgical or endoscopic simulation

With the tools described, surgical operations and stereotactic interventions or punctures can be simulated. Fig. 1.7 shows the way a craniotomy can be simulated and the access path to a region of interest explored

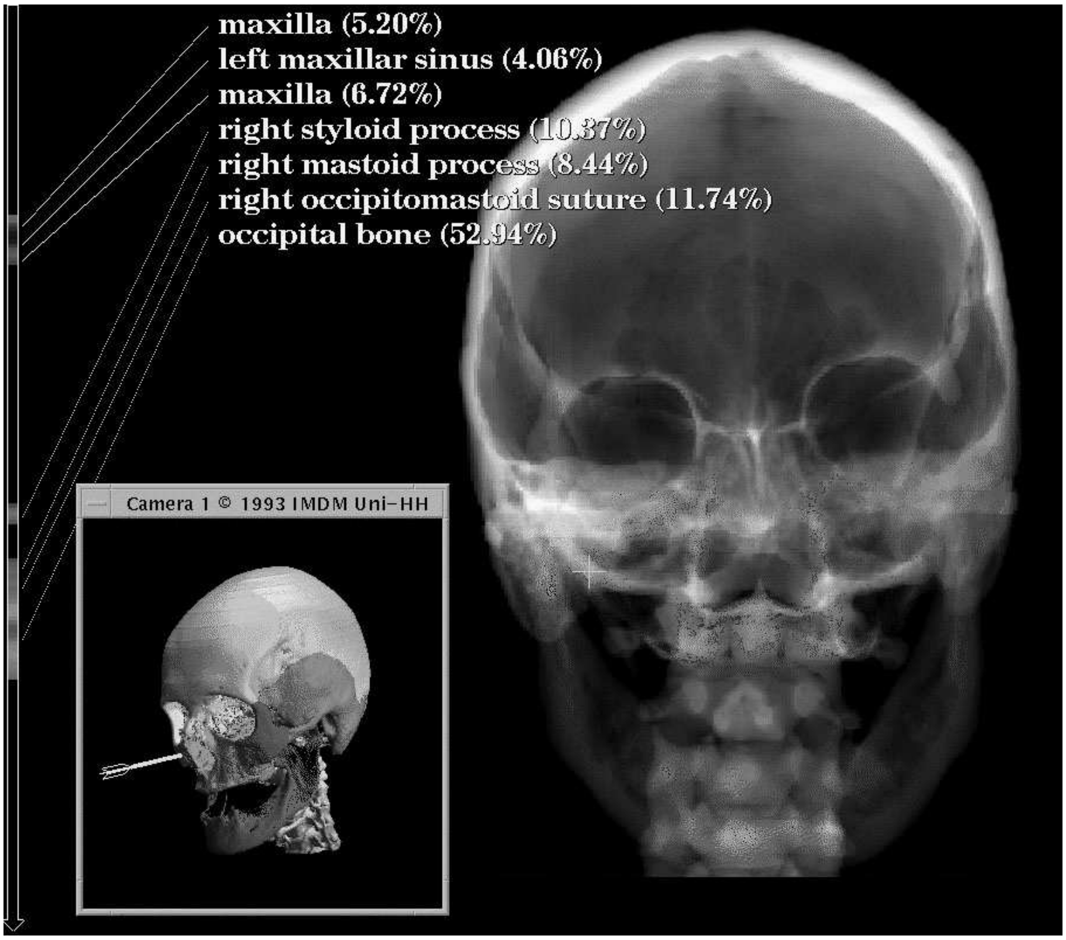

Figure 1.6: When the model contains Computer Tomographic data, X-ray images with any beam geometry and viewing direction can be simulated. At any location (cross) the contribution of different anatomic constituents to the total radiological density may be visualized.

[39]. Stereotactic surgery can be simulated as well. Objects encountered using a certain position and direction of the probe can be identified easily. The locations involved are expressed in the standard stereotactic coordinate system [55]. Endoscopic viewing can be simulated easily by simply placing the viewing position into the body and choosing a low focal length of the camera, if required even as binocular stereo view (Fig. 1.8).

1.4.4 Authoring system for other media

Beside the interactive use we can, of course, generate pictures for atlas books or multimedia systems. Also glas models can be simulated at the computer screen (Fig. 10) or even physical models can be created, for example with the technique of computer controlled stereo lithography. Or we can program scripts, that generate movies showing, for example, a dissection or the animation of the blood flow. As a novel feature, we do not have to edit the material again if the underlying knowledge changes. The same scripts (with possibly some modifications) would produce the new material, especially new pictures, automatically. A

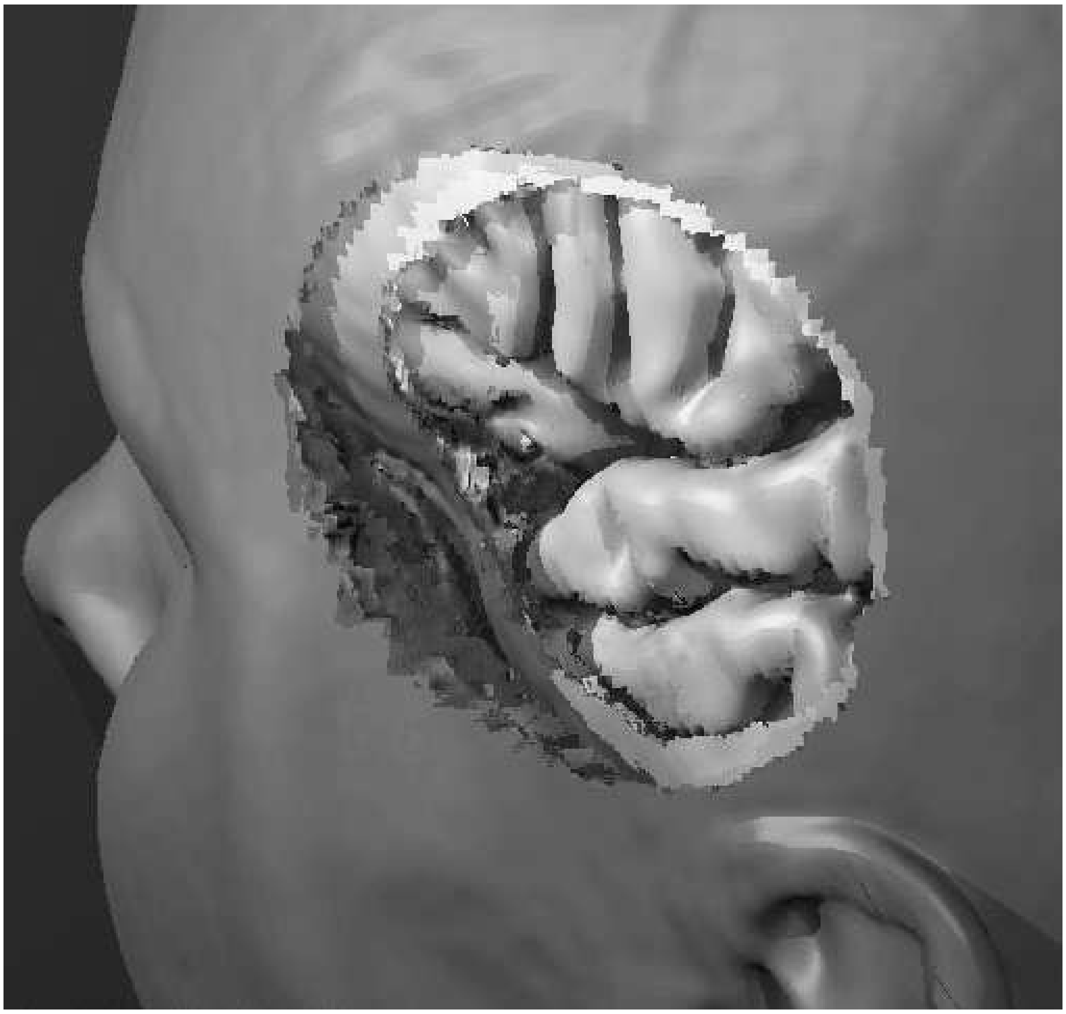

Figure 1.7: Rehearsal of surgical interventions. A “craniotomy” can be simulated. At any stage the user can inquire the knowledge base, e.g. about consequences of the intervention concerning brain function.

player for such “intelligent movies”, called VOXEL-MAN jr., is now available on CD-ROM [17].

1.4.5 Other uses

We are sure that the applications of the ‘intelligent volume’ approach, described above, are only a small selection of what is possible. Other physicians or researchers may have quite different ideas on how to use the knowledge represented in this form. This might be just as a pictorial dictionary (see Fig. 1.3) or as a spatial index for literature search. In the latter case one exotic application could be to color code the number of publications related to any part of the cerebral cortex.

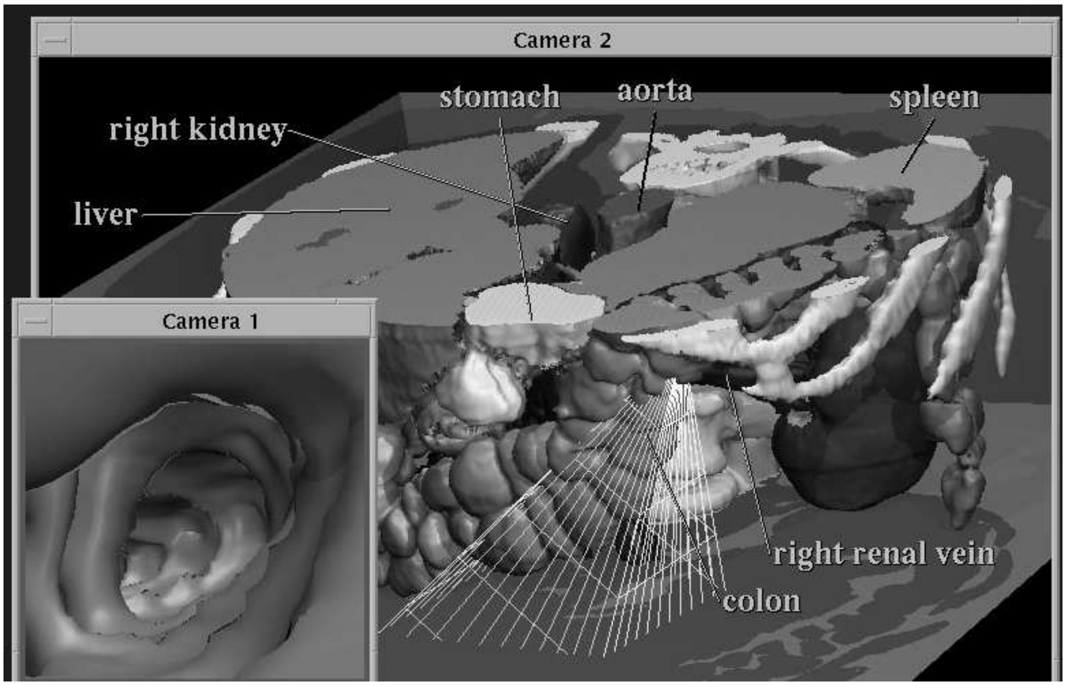

Figure 1.8: Endoscopic viewing within a 3D-atlas of the human abdomen. The field of view (shown as set of rays) may be chosen arbitrarily while the resulting image appears in a separate window (bottom left). Here a view into the colon is shown.

1.5 Perspectives

The new knowledge representation of human anatomy and function potentially exceeds all previous approaches concerning generality and versatility. It obviates the classical duality of text and pictures; pictures and text are not longer predefined but represent the users view to the knowledge representation. The user can literally compose his personal perspective both to the spatial and the symbolic content of the knowledge base. With the example of a 3D-atlas of the human brain and skull we have shown that the generality of the underlying model potentially allows us to derive all classical aids for visualizing anatomical and functional knowledge (cadaver dissection, pictures, classical atlases, movies) and to generate new ones not possible so far (simulation of surgery, medical imaging). One important new feature is that any application profits automatically from an update or refinement of the knowledge base.

While the feasibility of the general concept and the practical usefulness in the special application has been shown, a really general representation of the human body is still a long way off. One class of problems can be solved more ore less by hard work and more powerful computers: The spatial resolution has to be increased substantially by using higher resolution imaging including histological crossectional imaging [1], also on the microscopic level. However, this would for the time being exceed the storage capabilities of today's general purpose computers. Moreover, the model is static so far. It is imaginable to include a biomechanical model for simulation of motion and deformation (for animation of joints or heart motion or feeling mechanical properties using a virtual reality environment). Finite element techniques are available for this purpose, but still too complex to run interactively on an affordable computer.

Problems arising in general modelling are more severe. If we want to map standard textbook knowledge into a formal representation we find that there are inconsistencies arising either from historical development or from different views of the same situation. Even where this is not the case, the limits of semantic network modelling for our purpose are unclear. A further problem is the description of variations like interindividual or age differences. It is an extremely difficult task worked on at different places with no obvious solution [4, 5, 10]. The situation is similar in the case of "fuzzy" anatomical objects (e.g. the boundaries of functional cortical areas, tumors).

Although the present model is far from being general, its versatility already offers an virtually infinite number of ways for using it. So we are surprised at the different paths students navigate through the knowledge, and they often generate types of pictures which the authors have not thought of. The versatility is not without problems, because the user might easily get "lost in knowledge space". Therefore, as in classical book writing, an expert author is needed who limits and structures the knowledge according to his discipline and the audience. Yet instead of designing new drawings and writing text, he is structuring access to already stored knowledge by giving commands to a computer.

Acknowledgements

We are grateful for the invaluable incentives that Prof. Werner Lierse (late head of the Department of Neuroanatomy) provided in the initial phase of the project. Thanks go also to Dr. Rainer Maas (Dept. of Radiology), Dr. Uwe Rehder (Dept. of Orthopedics), Prof. Ernst Richter (Dept. of Pediatric Radiology), Prof. Hermann Zeumer (Dept. of Neuroradiology), and Siemens Medical Systems for providing tomographic data. We gratefully acknowledge the critical comments of Dr. John Morris (Dept. of Human Anatomy, University of Oxford). Last but not least we thank the members of our department for their enthusiastic support.