A new representation of knowledge concerning human anatomy and function

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

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

Probably no other subject has 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, as 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 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 organization 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 expected 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 engineering allows the structuring of symbolic information in a much more consistent and flexible way. Instead of dealing with static images or movies, computer graphics allow for the creation of realistic interactive three-dimensional (3D) models of human anatomy and function. Instead or in addition to a textual or formal description, dynamic 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 potentially to extract all classical knowledge representations for different purposes as views to this abstraction. More important, the word 'view' could be taken literally in that access to the information could be via a visual representation of the model.

It is clear that the generation of such an abstraction is a huge undertaking, and perhaps it will not be feasible. However, it may be possible to condense the spatial and symbolic description of the body in a hybrid data structure, which has the character of a model. An interactive 3D atlas of the human gross anatomy[1-3] based on such ideas has been made, so this concept is feasible.

Formalizing anatomical knowledge

The first efforts in developing digital representations of anatomical knowledge resulted in databases, 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[4]. 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 that can be found in standard books and atlases like the *Talairach-* or *Schaltenbrand-Atlas*[5,6]. Usually hyper- and multimedia techniques are used to link symbolic with pictorial data. Some of these systems include histological or photographic images from cadaver dissection. Representations of this concept are Neuro Database[7] or the database of the BrainMap project[8]. On the same basis, teaching programs have also been developed such as A.D.A.M.[9], or BrainStorm[10]. 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 that were designed to facilitate interpretation of images used for clinical diagnosis[11-15]. They have in common that anatomical structures are represented as sets of coordinates or 2D contours only, and the anatomical knowledge is restricted to a (sometimes hierarchical) list of anatomical names for a limited number of morphological objects. Likewise purely abstract descriptions without a pictorial representation have been described for automatic image interpretation with methods of computer visualization[16,17]. Only a few digital representations of anatomy based on a 3D model have been described[18-21]. These projects have in common objects represented as 3D contours only and are therefore hollow and thus do not correspond to reality.

A first step to acquiring detailed data for a space-filling model is proposed in the "Visible Human" project[22], 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 anatomic model.

'Intelligent volume' approach

Basic idea. The key idea underlying the new approach is to combine in one single framework a detailed spatial model enabling realistic visualization with a symbolic model of the human body (Fig. 1). The spatial model is sampled from a living subject by means of 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 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 with the image volume (Fig. 1). Further attribute volumes may be added that contain, for example, the incidence of a tumour type or a time tag for blood propagation on a per voxel basis.

The objects themselves bear attributes as well. They may be divided into two groups: attributes defining their visual appearance, like colour, texture and 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 for certain features such as 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 artificial intelligence[23] the technique of semantic networks was chosen[17,24]. This allows the expression of links between objects via relations with given semantics. Three groups have been considered so far: those for modelling abstractions, structure and functions. The following examples outline the essentials of the semantic network model (see ref. 25 for details).

Modelling descriptive information. A small section of the network containing some relations concerning the right optical nerve is shown in Fig. 1. 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 according to morphology, the right optical nerve is a part of the brain, whereas in functional anatomy it is part of the visual system.

Functional relations such as stimulus propagation may be modelled via the **Projecting To** 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 **Supplied By** relation. Finally, the abstraction saying that it belongs to the class of cranial tracts is expressed by the **Is A** 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 and/or areas are involved in stimulus propagation, and beyond that, we could enquire about consequences of a failure of the right ophthalmic artery.

Despite the elegance of the approach it must be clear that today's medical knowledge is not sufficiently consistent to allow a universally agreed-upon description. Thus only part of the available knowledge can be represented in

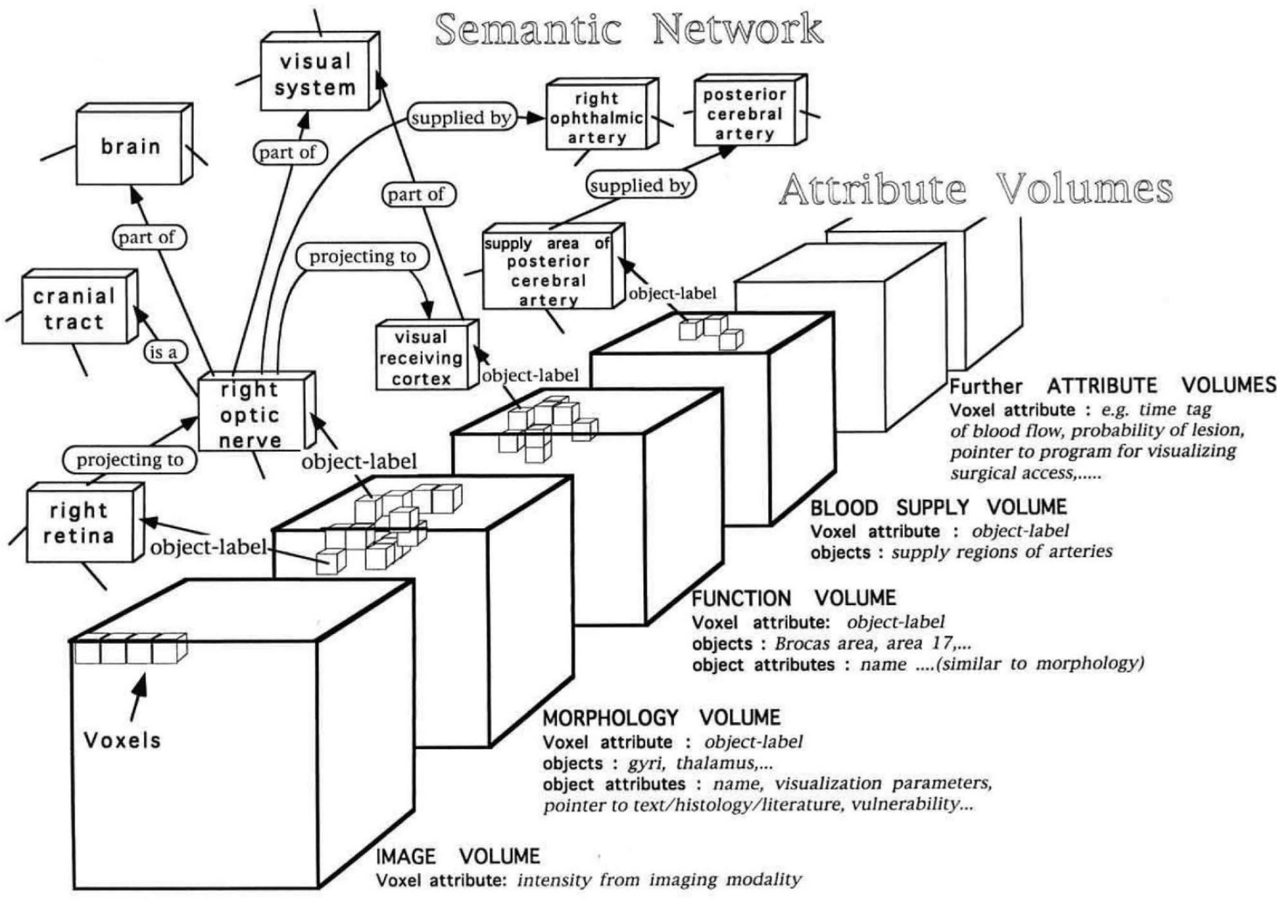

Fig. 1 Block diagram of the knowledge representation consisting of a semantic network linked to a digital volume representation.

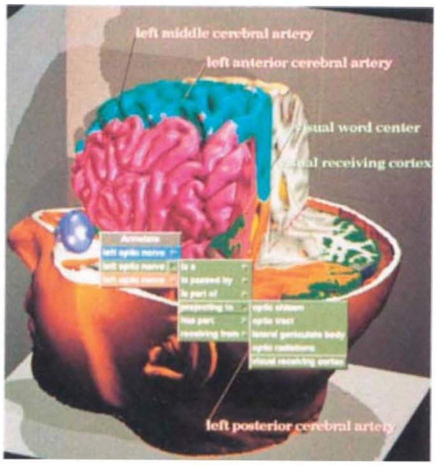

Fig. 2 Exploration of the brain. Arbitrary cutting reveals the interior. The user has gained access to the information available on the optic nerve concerning functional anatomy, which appears as a cascade of pop-up menus. He has asked the system to colour-mark some blood supply areas and cortical areas. He can derive from the colours that the visual word center is supplied by the left middle cerebral artery.

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 are, however, linked to the respective objects in the semantic network via pointers (see Fig. 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 continuous process driven by practical experience with its application.

Anatomic model of the VOXEL-MAN. With the tools described above, a semantic network model of the body has been built, which contains, so far, 1,000 objects with 2,500 links in the domains of morphology, pathology, functional anatomy and the vascular system. It is being refined continuously. It is called the generic knowledge base, because it contains general knowledge that is independent of an actual spatial visual representation (such as 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 loses 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 by magnetic resonance tomography with resolution of 1.5 mm, the skull was sampled by 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[26,27]. 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 3D 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 visualization parameters (colour, reflectivity and so on) of 'cortex'. Even if advanced editing tools are used, the process of generating an intelligent volume is tedious. For the atlas of the skull and brain it took on the order of one man-year to segment the 300 anatomic objects available so far.

Fig. 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. Access to any voxel may be achieved in several languages (bottom left).

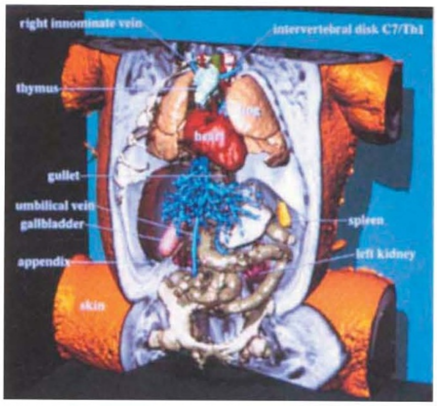

Fig. 4 View generated with the interactive 3D atlas of a human fetus, derived from magnetic resonance imaging.

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 provide all effects that a painter would use to paint a 3D scene onto a planar screen just by applying simple laws of projective geometry and optics[28,29]. In the resulting system, a viewer can generate arbitrary perspective views by choosing a certain viewpoint, 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 and/or remove objects, to

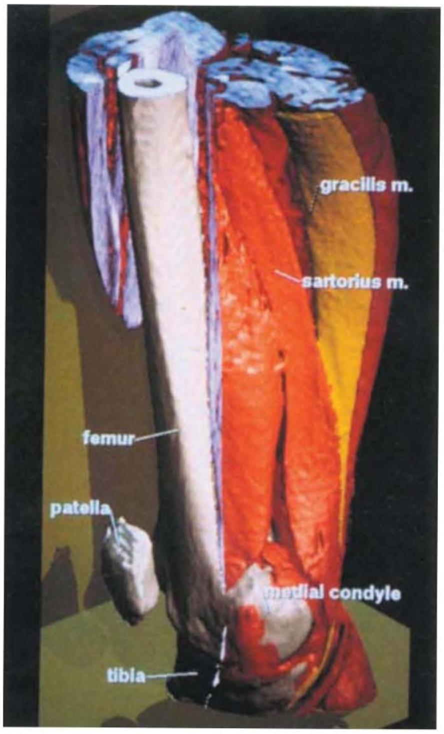

Fig. 5 View generated with the 3D interactive atlas of the musculoskeletal system derived from computed tomography.

colour-mark related objects and to show corresponding histological sections. It also provides information as subnets describing relations of the object (as seen in Fig. 2) or a textual description. The VOXEL-MAN system is presently being validated in 25 institutions around the world. For more details of the implementation we refer to ref. 2 .

Intelligent volumes — possibilities

The result of the use of intelligent volumes 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.

Learning anatomy and function. In the case of a 3D atlas of the brain[30], 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. 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 coloured as well. The language for the annotations can be chosen from English, French, German, Japanese and Latin (see Fig. 3). In the case of the 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 (for example, the lobe to which a gyrus belongs) or functional regions traversed by a stimulus can be queried and coloured, displayed, removed, and so on (Fig. 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 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 and then point to the included regions to obtain the colour-marked functional regions (Fig. 2).

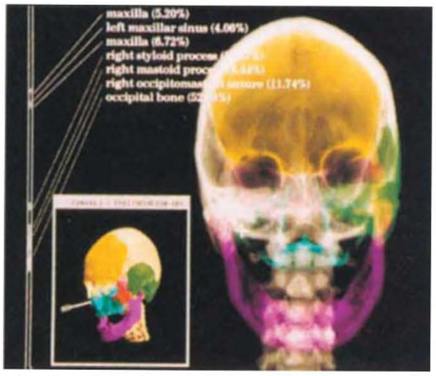

Fig. 6 When the model contains computed tomographic data, X-ray images with any beam geometry and viewing direction can be simulated. At any location (cross) the contribution of different anatomic constituents to the total radiological density may be visualized.

Fig. 7 Rehearsal of surgical interventions. A 'craniotomy' can be simulated. At any stage the user can inquire the knowledge base, for example, about consequences of the intervention concerning brain function.

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 neighbouring objects as shown in Fig. 5. In this case, it is advisable to add the ventricular system for orientation, which, unlike that seen 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. 4) and the musculoskeletal system (Fig. 5).

Radiological image interpretation. Because the model is derived from radiological imagery, it lends itself to teaching and reference in radiology. As shown in Fig. 6, this can be done by presenting cross-sectional images in the context of basic anatomy. 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. Because 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. 6).

Surgical or endoscopic simulation. With the tools described, surgical operations and stereotactic interventions or punctures can be simulated. Fig. 7 shows the way a craniotomy can be simulated and the access path to a region of interest explored. Stereotactic surgery can be simulated as well. Objects encountered using a certain position and direction of the probe can be

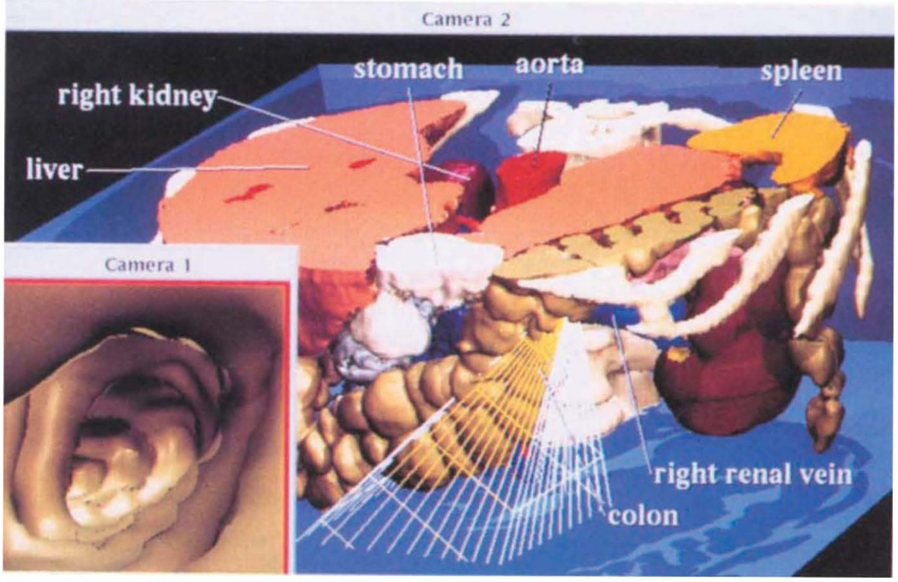

Fig. 8 Endoscopic viewing within a 3D-atlas of the human abdomen. The field of view (shown as set of rays) may be chosen arbitrarily while the resulting image appears in a separate window (bottom left). Here a view into the colon is shown.

identified easily. The locations involved are expressed in the standard stereotactic coordinate system[5]. Endoscopic viewing can be simulated by simply placing the viewing position into the body and choosing a low focal length of the camera (Fig. 8). If required, a binocular stereo view can be selected.

Authoring system for other media. Besides the interactive use we can, of course, generate pictures for atlas books or multimedia systems. Also glass models can be simulated at the computer screen (see illustration in table of contents) 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.

Other uses. We are sure that the applications of the intelligent volume approach, described above, are only part 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. 3) or as a spatial index for literature search. In the latter case one exotic application could be to colour code the number of publications related to any part of the cerebral cortex.

Perspectives

The new knowledge representation of human anatomy and function potentially exceeds all previous approaches in generalizability and versatility. It obviates the classical duality of text and pictures; pictures and text are no longer predefined but represent the user's view of the knowledge representation. The user can literally compose his personal perspective both on 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 previously impossible (simulation of surgery, medical imaging). One important new feature is that any application profits automatically from an update or refinement of the knowledge base.

Although the feasibility of the general concept and the practical usefulness in the special application have been shown, a really general representation of the human body is still a long way off. One class of problems can be solved more or less by hard work and more powerful computers: The spatial resolution has to be increased substantially by using higher resolution imaging including histological cross-sectional imaging[22], 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. This is an extremely difficult task being tackled by several groups but with no obvious solution[31–33]. The situation is similar in the case of 'fuzzy' anatomical objects (for example, the boundaries of functional cortical areas or tumours).

Although the present model is far from being general, its versatility already offers an infinite number of ways for using it. We are surprised at the different paths students use to navigate through the knowledge, often generating types of pictures that the authors had 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.

Institut für Mathematik und
Datenverarbeitung in der Medizin
Universität Hamburg
Martinistraße 52
20246 Hamburg, Germany

Acknowledgements

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

1. Höhne, K.H. *et al.* A 3D anatomical atlas based on a volume model. in *IEEE Comput. Graphics Appl.* **12**, 72–78 (1992).
2. Tiede, U. *et al.* A computerized three-dimensional atlas of the human skull and brain. *Am. J. Neuroradiol.* **14**, 551–559 (1993).
3. Schubert, R. *et al.* A new method for practicing exploration, dissection, and simulation with a complete computerized three-dimensional model of the brain and skull. *Acta Anat.* **150**, 69–74 (1994).
4. Bloom, F.E. Databases of brain information. in *Three-Dimensional Neuroimaging* (ed., A. W. Toga), Ch. 13, 273–306 (Raven, New York, 1990).
5. Talairach, J. & Tournoux, P. *Co-planar Stereotaxic Atlas of the Human Brain* (Thieme, Stuttgart, 1988).
6. Schaltenbrand, G. & Wahren, W. *Atlas for Stereotaxy of the Human Brain* (Thieme, Stuttgart, 1977).
7. Wertheim, S.L. The brain database: A multimedia database for neuroscience research and teaching. in *13th Annual Symposium on Computer Applications in Medical Care* (eds Lemke, H.U., Rhodes, M.L., Jaffe, C.C. & Felix, R.) 399–404 (IEEE Computer Society Press, Los Alamitos, California, 1989).
8. Fox, P.T., Mikiten S., Davis, G. & Lancaster, J.L. BrainMap: A database of human functional brainmapping. in *Functional Neuroimaging: Technical Foundations* (eds Thatcher, R.W., Hallett, M., Zeffiro, T. & John, E.R.), ch. 9, 95–105 (Academic, Orlando, Florida, 1994).
9. Lipscomb, K. & Kittrel, A. Software gets under your skin. *Physicians Comput.* **9**, 14–16 (1991).
10. Dev, P., Coppa, G.P. & Tancred, E. Brainstorm: Designing an interactive neuroanatomy atlas [abstract]. *Radiology* **185**(P), 413 (1992).
11. Evans, A.C. *et al.* Anatomical-functional correlation using an adjustable MRI-based region of interest atlas with positron emission tomography. *J. Cereb. Blood Flow Metab.* **8**, 513–530 (1988).
12. Greitz, T., Bohm, C., Holte, S. & Eriksson, L. A computerized brain atlas: Construction, anatomical content, and some applications. *J. Comput. Assist. Tomogr.* **15**, 26–38, (1991).
13. Keyserlingk, D.G., Niemann, K., Wasel, J. & Maurer, K. Digital brain atlas for structure localization in CT combined with EEG mapping. *Psychiat. Res.* **29**, 461–462 (1989).
14. Lehmann, E.D.M. *et al.* Computer-aided interpretation of SPECT images of the brain using an MRI-derived 3D neuro-anatomical atlas. *Med. Inform.* **16**, 151–166 (1991).
15. Niemann, K., Keyserlingk, D.G.v. & Wasel, J. Superimposition of an averaged three-dimensional pattern of brain structures on CT scans. *Acta Neurochir.* **93**, 61–67 (1988).
16. Carlsen, I.C. *et al.* Knowledge based interpretation of cranial MR images. in *Computer Assisted Radiology, Proc. CAR '91* (eds Lemke, H.U., Rhodes, M.L., Jaffe, C.C. & Felix, R.) 277–282 (Springer, Berlin, 1991).
17. Robinson, G.P., Colchester, A.C.F. & Griffin, L.D. Model-based recognition of anatom ical objects from medical images. *Image Vis. Comput.* **12**, 499–507 (1994).
18. Mano, I., Suto, Y, Suzuki, M. & Iio, M. Computerized three-dimensional normal atlas. *Radiat. Med.* **8**, 50–54 (1990).
19. Brinkley, J.F., Eno, K. & Sundsten, J.W. Knowledge-based client-server approach to structural information retrieval: The Digital Anatomist Browser. *Comput. Methods Programs Biomed.* **40**, 131–145 (1993).
20. Wahler-Lück, M., Schütz, T. & Kretschmann, H.-J. A new anatomical representation of the human visual pathways. *Graefe's Arch. Clin. Exp. Ophthalmol.* **229**, 201–205 (1991).
21. Gee, J.C., Reivich, M. & Bajcsy, R. Elastically deforming a 3D atlas to match anatomical brain images. *J. Comput. Assist. Tomogr.* **17**, 225–236 (1993).
22. Ackerman, M.I. Viewpoint: The Visible Human project. *J. Biocommun.* **18**, 14 (1991).
23. Winston, P.H. *Artificial Intelligence* 3rd edn (Addison-Wesley, Reading, Massachusetts, 1992).
24. Niemann, H., Sagerer, G.F., Schröder, S. & Kummert, F. Ernest: A semantic network system for pattern understanding. *IEEE Trans. Pattern Analyt. Machine Intell.* **12**, 883–905 (1990).
25. Pommert, A. *et al.* Symbolic modeling of human anatomy for visualization and simulation. in *Visualization in Biomedical Computing* (1994), Proc. SPIE 2359 (ed. Robb, R.A.) 412–423 (Rochester, Minnesota, 1994).
26. Höhne, K.H. & Hanson, W.A. Interactive 3D-segmentation of MRI and CT volumes using morphological operations. *J. Comput. Assist. Tomogr.* **16**, 285–294 (1992).
27. Schiemann, Th., Bomans, M., Tiede, U. & Höhne, K.H. Interactive 3D-segmentation in *Visualization in Biomedical Computing II*, Proc. SPIE 1808 (ed. Robb, R.A.) 376–383 (Chapel Hill, North Carolina, 1992).
28. Höhne, K.H. *et al.* 3D-visualization of tomographic volume data using the generalized voxel-model. *Visual Comput.* **6**, 28–36 (1990).
29. Tiede, U. *et al.* Investigation of medical 3D-rendering algorithms. *IEEE Comput. Graphics Appl.* **10**, 41–53 (1990).
30. Voxel-man brain and skull, an interactive 3D atlas for teaching and studying anatomy and its radiological appearance. (Springer Electronic Media, Heidelberg & New York, 1995).
31. Bookstein, F. *Morphometric Tools for Landmark Data.* (Cambridge Univ. Press, New York, 1991).
32. Brechbühler, C., Gerig, G. & Kübler, O. Surface parametrisation and shape description. in *Proc. First Conference on Visualization in Biomedical Computing, VBC '92*, 80–89 (Chapel Hill, North Carolina, 1992).
33. Declerck, J., Subsol, G., Thirion, I. & Ayache, N. Automatic retrieval of anatomical structures in 3d medical images. in *Proc. First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine* 153–162 (CVR Med, Nice, 1995).