

Advanced Applications of Volume Visualization Methods in Medicine

Georgios Sakas
Fraunhofer Institute for Computer Graphics
Wilhelminenstr. 7
64283 Darmstadt, Germany
gsakas@igd.fhg.de

Andreas Pommert
Institute of Mathematics and Computer Science in Medicine (IMDM)
University Hospital Eppendorf
Martinistraße 52
20251 Hamburg, Germany
pommert@uke.uni-hamburg.de

Abstract

Tomographic medical imaging techniques have become more popular in recent years. The wide availability of CT, MRI and Ultrasound in most large hospitals results in a rapidly increasing number of examinations with these devices. The State of The Art Report summarises the application of techniques developed over the recent years for visualising volumetric medical data common in modern medical imaging modalities such as CT, MRA, MRI, Nuclear Medicine, 3D-Ultrasound, Laser Confocal Microscopy etc. Although all of the modalities listed above provide "slices of the body", significant differences exist between the image content of each modality. The focus of the Report is be less in explaining algorithms and rendering techniques, but rather to point out their applicability, benefits, and potential in the medical environment.

In the first part, methods for all steps of the volume visualization pipeline from data preprocessing to object display are reviewed, with special emphasis on data structures, segmentation, and surface- and volume-based rendering. Furthermore, multimodality matching, intervention rehearsal, and aspects of image quality are discussed.

In the second part applications are illustrated from the areas of craniofacial surgery, traumatology, neurosurgery, radiotherapy, and medical education. Furtherly, some new applications of volumetric methods are presented: 3D ultrasound, laser confocal datasets, and 3D-reconstruction of cardiological datasets, i.e. vessels as well as ventricles. These new volumetric methods are currently under development but due to their enormous application potential they are expected to be clinically accepted within the next years.

Part I

Volume Visualization Principles

1 Introduction

Medical imaging technology has experienced a dramatic change over the past two decades. Previously, only X-ray radiographs were available which showed the depicted organs as superimposed shadows on photographic film. With the advent of modern computers, new *tomographic* imaging modalities like computed tomography (CT), magnetic resonance imaging (MRI), and positron emission tomography (PET) could be developed which deliver cross-sectional images of a patient's anatomy and physiology. These images show different organs free from overlays with unprecedented precision. Even the three-dimensional (3D) structure of organs can be recorded if a sequence of parallel cross-sections is taken.

For many clinical tasks like surgical planning, it is necessary to understand and communicate complex and often malformed 3D structures. Experience has shown that the "mental reconstruction" of objects from cross-sectional images is extremely difficult and strongly depends on the observer's training and imagination. For these cases, it is certainly desirable to present the human body as a surgeon or anatomist would see it.

The aim of *volume visualization* (also known as *3D imaging*) in medicine is to create precise and realistic views of objects from medical volume data. The resulting images, even though they are of course two-dimensional, are often called *3D images* or *3D reconstructions* to distinguish them from 2D cross-sections or conventional radiographs. The first attempts date back to the late 1970s, with the first clinical applications reported on the visualization of bone from CT in craniofacial surgery and orthopedics. Methods and applications have since been extended to other subjects and imaging modalities. The same principles are also applied to sampled and simulated data from other domains, such as fluid dynamics, geology, and meteorology [44].

2 Methods

An overview of the volume visualization pipeline as presented here is shown in fig. 1. After the acquisition of a series of tomographic images of a patient, the data usually undergoes some preprocessing for data conversion and possibly image filtering. From this point, one of several paths may be followed.

The dotted line in fig. 1 represents an early approach where an object is reconstructed from its contours on the cross-sectional images. All other methods, represented by the solid line, start from a contiguous *data volume*. If required, equal spacing in all three directions can be achieved by interpolation. Like a 2D image, a 3D volume can be filtered to improve image quality. Corresponding to the *pixels* (picture elements) of a 2D image, volume elements are called *voxels* (volume elements).

The next step is to identify the different objects represented in the data volume so that they can be removed or selected for visualization. The simplest way is to binarize the data with an intensity threshold, e.g. to distinguish bone from other tissues in CT. Especially for MRI data, however, more sophisticated *segmentation* methods are required.

After segmentation, there is a choice which *rendering* technique is to be used. The more traditional surface-based methods first create an intermediate surface representation of the object to be shown. It may then be rendered with any standard computer graphics method. More recently, volume-based methods have been developed which create a 3D view directly from the volume data. These methods use the full gray level information to render surfaces, cuts, or transparent and semi-transparent volumes. As a third way, transform-based rendering methods may be used.

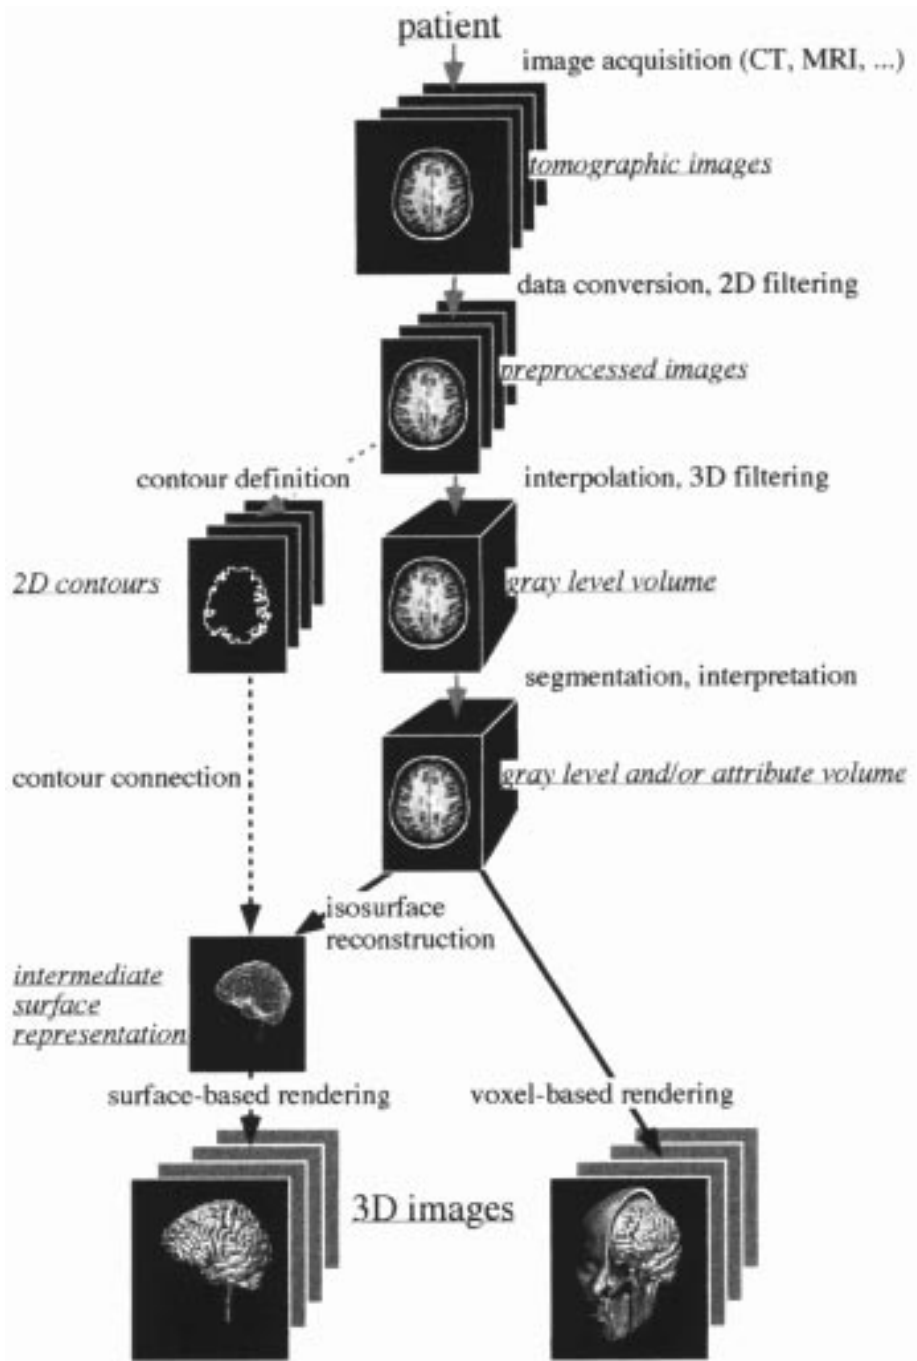


Figure 1: Overview of the volume visualization pipeline. Individual processing steps may be left out, combined, or reversed in order by a particular method.

Extensions to the volume visualization pipeline not shown in fig. 1 but also covered here include multimodality matching and intervention rehearsal.

2.1 Preprocessing

The data we consider usually comes as a spatial sequence of 2D cross-sectional images. If they are put on top of each other, a contiguous *gray level volume* is obtained. The resulting data structure is an orthogonal 3D array of voxels, each representing an intensity value. This is called the *voxel-model*.

Many algorithms for volume visualization work on *isotropic* volumes where the sampling density is equal in all three dimensions. In practice, however, only very few data sets have this property, especially for CT. In these cases, the missing information has to be reconstructed in an *interpolation* step. A quite simple method is linear interpolation of the intensities between adjacent images. Higher order functions such as splines usually give better results for fine details [59]. Shape-based methods are claimed to be superior in certain situations [7]; however these are depending on the results of a previous segmentation step.

With respect to later processing steps such as segmentation, it is often desirable to improve the signal-to-noise ratio of the data, using image or volume filtering. Well-known *noise filters* are average, median and Gaussian filters [75]. These methods, however, tend to smooth out small details as well; better results are obtained with *anisotropic diffusion* filters which largely preserve object boundaries [28].

2.1.1 Data Structures for Volume Data

There are a number of different data structures for volume data. The most important are

- *binary voxel-model*: voxel values are either 1 (object) or 0 (no object). This very simple model is not much in use any more. In order to reduce storage requirements, binary volumes may be subdivided recursively into subvolumes of equal value; the resulting data structure is called an *octree*.
- *gray level voxel-model*: each voxel holds an intensity information. Octree representations have also been developed for gray level volumes [50].
- *generalized voxel-model*: in addition to an intensity information, each voxel contains *attributes*, describing its membership to various objects, and/or data from other sources (e.g. MRI and PET) [38].
- *intelligent volumes*: as an extension of the generalized voxel model, properties of objects (such as color, names in various languages, pointers to related information) and their relationships are modeled on a symbolic level [40, 72, 88]. This data structure is the basis for advanced applications such as medical atlases (see below).

2.2 Segmentation

A gray level volume usually represents a large number of different structures obscuring each other. To display a particular one, we thus have to decide which parts of the data we want to use or ignore. Ideally, selection would be done with a command like “show only the brain”. This, however, requires that the computer knows which parts of the volume (or, more precisely, which voxels) constitute the brain and which do not.

A first step towards object recognition is to partition the gray level volume into different regions which are homogeneous with respect to some formal criteria and corresponding to real (anatomical) objects. This process is called *segmentation*. The generalized voxel-model is a suitable data structure for representing the results. In a further *interpretation* step, the regions may be identified and labeled with meaningful terms such as “white matter” or “ventricle”.

All segmentation methods can be characterized as being either “binary” or “fuzzy”, corresponding to the principles of binary and fuzzy logic, respectively [108]. In *binary segmentation*, the question whether a voxel

belongs to a certain region is always answered yes or no. This information is a prerequisite e.g. for creating surface representations from volume data. As a drawback, however, uncertainty or cases where an object takes up only a fraction of a voxel (*partial volume effect*) cannot be handled properly. For example, a very thin bone would appear with false holes on a 3D image. Strict yes-no decisions are avoided in *fuzzy segmentation*, where a set of probabilities is assigned to every voxel, indicating the evidence for different materials. Fuzzy segmentation is closely related to the so-called volume rendering methods discussed later.

Currently, a large number of segmentation methods for 3D medical images are being developed, which may be roughly divided into three classes: point-, edge-, and region-based methods. The methods described often have been tested successfully on a number of cases; experience has shown, however, that the results should always be used with care.

2.2.1 Point-Based Segmentation

In point-based segmentation, a voxel is *classified* only depending on its intensity, no matter where it is located. A very simple but nevertheless important example which is very much used in practice is *thresholding*: a certain intensity range is specified with lower and upper threshold values. A voxel belongs to the selected class if and only if its intensity level is within the specified range. Thresholding is the method of choice for selecting bone or soft tissue in CT. In volume-based rendering, it is often performed during the rendering process itself so that no explicit segmentation step is required.

In order to avoid the problems of binary segmentation, Drebin et al. use a fuzzy maximum likelihood classifier which estimates the percentages of the different materials represented in a voxel, according to Bayes' rule [21]. This method requires that the gray level distributions of different materials are different from each other and known a-priori. This is approximately the case in musculoskeletal CT.

Unfortunately, these simple segmentation methods are not suitable if different structures have mostly overlapping or even identical gray level ranges. This situation frequently occurs e.g. in the case of soft tissues from CT or MRI. The situation is somewhat simplified if multiple-parameter data are available, such as T_1 - and T_2 -weighted images in MRI, emphasizing fat and water, respectively. In this case, individual threshold values can be specified for every parameter. To somewhat generalize this concept, voxels in an n-parameter data set can be considered as n-dimensional vectors in an n-dimensional *feature space*. In *pattern recognition*, this feature space is partitioned into subspaces, representing different tissue classes or organs. This is called the *training phase*: in supervised training, the partition is derived from feature vectors which are known to represent particular tissues [16, 29]. In unsupervised training, the partition is automatically generated [29]. In the subsequent *test phase*, a voxel is classified, according to the position of its feature vector in the partitioned feature space.

With especially adapted image acquisition procedures, pattern recognition methods have successfully been applied to considerable numbers of two- or three-parametric MRI data volumes [16, 29]. Quite frequently, however, isolated voxels or small regions are incorrectly classified (e.g. subcutaneous fat in the same class as white matter). To eliminate these errors, a connected component analysis (see below) is often applied.

A closely related method is based on *neural network* methodology, as developed by Kohonen [48]. Instead of an n-dimensional feature space, a so-called *topological map* of $m \times m$ n-dimensional vectors is used. During the training phase, the map iteratively adapts itself to a set of training vectors which may either represent selected tissues (supervised learning) or the whole data volume (unsupervised learning) [31, 99]. Finally, the map develops several relatively homogeneous regions, which correspond to different tissues or organs in the original data. The practical value of the topological map for 3D MRI data seems to be generally equivalent to that of pattern recognition methods.

2.2.2 Edge-Based Segmentation

The aim of edge-based segmentation methods is to detect intensity discontinuities in a gray level volume. These edges (in 3D, they are actually surfaces; it is however common to speak about edges) are assumed to represent the borders between different organs or tissues. Regions are subsequently defined as the enclosed areas.

A common strategy for edge detection is to locate the maxima of the first derivative of the 3D intensity function. A method which very accurately locates the edges was developed by Canny [14]. All algorithms using the first derivative, however, share the drawback that the detected contours are usually not closed, i.e. they do not separate different regions properly.

An alternative approach is to detect zero-crossings of the second derivative. The Marr-Hildreth operator convolves the input data with the Laplacian of a Gaussian; the resulting contour volume describes the locations of the edges [58]. With a 3D extension of this operator, Bomans et al. segmented and visualized the complete human brain from MRI for the first time [12]. Occasionally, however, this operator creates erroneous “bridges” between different materials which have to be removed interactively. Also, location of the surfaces is not always satisfactory.

Snakes [43] are 2D image curves, which are adjusted from an initial approximation to image features by a movement of the curve caused by simulated forces. The so called external force is produced by image features. An internal tension of the curve resists against highly angled curvatures, which makes the Snakes movement robust against noise. After a starting position is given it is adapted to an image by relaxation to an equilibrium of the external force and internal tension. To calculate the forces an external energy has to be defined. The gradient of this energy is proportional to the external force. Defining the external energy as the distance to the next maximum of the opacity function [52] produced good results.

The segmentation by Snakes is due to its 2D definition performed in a slice-by-slice manner. I. e. the resulting curves for a slice are copied into the neighboring slice and the minimization is started again. The segmentation process may be controlled by the user, by stopping the automatic tracking, if the curves run out of the contours and define a new initial curve.




Figure 2: The principle of Segmentation by Snakes

For this reason two methods have been applied to enter an initial curve for the Snake. The first is the interactive input of a polygon. Since the Snake contracts due to its internal energy, the contour to be segmented has to be surrounded by this polygon. The second one is a contour tracing method, using an A^* search tree to find the path with minimal costs between two interactively marked points [100], [65].

The quality of the result depends on the similarity of two adjacent slices. Normally, this is varying within a dataset. Therefore, in regions with low similarity, the slices to be segmented by the interactive method must be selected rather tightly.

2.2.3 Region-Based Segmentation

Region-based segmentation methods consider whole regions instead of individual voxels or edges. Since we are actually interested in regions, this approach appears to be the most natural. Properties of a region are e.g. its size, shape, location, variance of gray levels, and its spatial relation to other regions.

A typical application of region-based methods is to post-process the results of a previous point-based segmentation step. For example, a *connected component analysis* may be used to determine whether the voxels which have been classified as belonging to the same class are part of the same (connected) region. If not, there are likely errors in classifying the voxels in the smaller regions.

A practical interactive segmentation system based on the methods of *mathematical morphology* was developed by Höhne and Hanson [39, 82]. Regions are initially defined with thresholds; the user can subsequently apply simple but fast operations such as *erosion* (to remove small “bridges” between erroneously connected parts), *dilation* (to close small gaps), connected components analysis, region fill, or Boolean set operations. Segmentation results are immediately visualized on orthogonal cross-sections and 3D images, such that they may be corrected or further refined in the next step (fig. 3). With this system, segmentation of gross structures is usually a matter of minutes.

For automatic segmentation, the required knowledge about data and anatomy needs to be represented in a suitable model. A comparatively simple approach is presented by Brummer et al., who use a fixed sequence of morphological operations for the segmentation of brain from MRI [13]. For the same application, Raya and Udupa developed a rule-based system which successively generates a set of threshold values [74]. Rules are applied depending on measured properties of the resulting regions. Bomans generates a set of object hypotheses for every voxel, depending on its gray level [11]. Location, surface-volume ratio etc. of the resulting regions are compared to some predefined values, and the regions are modified accordingly. Menhardt uses a rule-based system which models the anatomy with relations such as “brain is inside skull” [60]. Regions are defined as fuzzy subsets of the volume, and the segmentation process is based on fuzzy logic and fuzzy topology.

One of the problems of these and similar methods for automatic segmentation is that the required anatomical knowledge is often represented in more or less ad-hoc algorithms, rules, and parameters. A more promising approach is to use an explicit three-dimensional organ model. For the brain, Arata et al. developed an atlas of the “normal” anatomy and its variation in terms of a probabilistic spatial distribution, obtained from 22 MRI data sets of living persons [3]. The model was reported suitable for the automatic segmentation of various brain structures, including white matter lesions. A similar approach is described in [47].

Another interesting idea is to investigate object features in *scale-space*, i.e. at different levels of image resolution. This approach allows to ignore irrelevant image detail. One such method developed by Pizer et al. considers the symmetry of previously determined shapes, described by medial axes [62]. The resulting ridge function in scale-space is called the *core* of an object. It may be used e.g. for interactive segmentation, where the user can select, add or subtract regions, or move to larger “parent” or smaller “child” regions in the hierarchy. Other applications like automatic segmentation or registration are currently being investigated.




Figure 3: 3D image of a still-born fetus from MRI. About 50 objects were defined using an interactive segmentation system.

In conclusion, automatic segmentation systems are not yet robust enough to be generally applicable to medical volume data. Interactive segmentation which combines fast operations with the unsurpassed human recognition capabilities is still the most practical approach.

2.3 Surface-Based Rendering

The key idea of *surface-based* rendering methods is to extract an intermediate surface description of the relevant objects from the volume data. Only this information is then used for rendering. If triangles are used as surface elements, this process is called *triangulation*.

A clear advantage of surface-based methods is the possibly very high data reduction from volume to surface representations. Resulting computing times can be further reduced if standard data structures such as triangle meshes are used which are supported by common rendering hard- and software.

On the other hand, the *surface reconstruction* step throws away most of the valuable information on the cross-sectional images. Even simple cuts are meaningless because there is no information about the interior of an object. Furthermore, every change of surface definition criteria such as thresholds requires a recalculation of the whole data structure.

The first method to be widely used in clinical practice, known as the *cube model*, was developed by Herman et al. [15]. The gray level volume is first binarized with an intensity threshold. Then, a list of square voxel faces is created which denote the border between voxels in- and outside the object.

A more recent method by Lorensen and Cline called *marching cubes* creates an *isosurface*, representing the locations of a certain intensity value in the data [56]. This algorithm basically considers a cube of $2 \times 2 \times 2$ contiguous voxels. Depending on whether one or more of these voxels are inside the object (i.e. above a

threshold value), a surface representation of up to four triangles is placed within the cube. The exact location of the triangles is found by linear interpolation of the intensities at the voxel vertices. The result is a highly detailed surface representation with subvoxel resolution (fig. 4). Surface orientations are calculated from gray level gradients. Meanwhile, a whole family of similar algorithms has been developed [68, 102, 106].




Figure 4: Triangulated (top) and shaded (bottom) portion of the brain from MRI, created with the marching cubes algorithm.

Applied to clinical data, the marching cubes algorithm typically creates hundreds of thousands of triangles. As has been shown, these numbers can be reduced considerably by a subsequent simplification of the triangle meshes, without much loss of information [87, 107].

An even simpler approach to surface reconstruction developed by the same group uses points instead of triangles [17]. This method, called *dividing cubes*, subdivides a group of $2 \times 2 \times 2$ contiguous voxels into smaller subcubes, whereby the intensities are interpolated. The surface description is made from those subcubes that approximate the threshold value. Similar to the marching cubes algorithm, every subcube contains a surface normal, calculated from gray level gradients.

A different method for the reconstruction of the polygonal mesh from the segmentation results is based on the Delauney interpolation developed and published by Boissonnat [10]. Using this method the volume of the contours is computed by a three-dimensional triangulation which allows an extraction of the surface of the object. An extreme complexity of the surface model can be avoided and, at the same time, a high approximation quality can be achieved, but to guarantee a real-time interaction with the anatomic model an additional reduction step is necessary. Fig. 5 shows the result of the triangle reduced surface of the Virtual Human Project patient. The reduction method can be parameterized and thus allows to derive models of different levels of detail. Even high reduction rates do not imply loss of structure of the anatomic object.

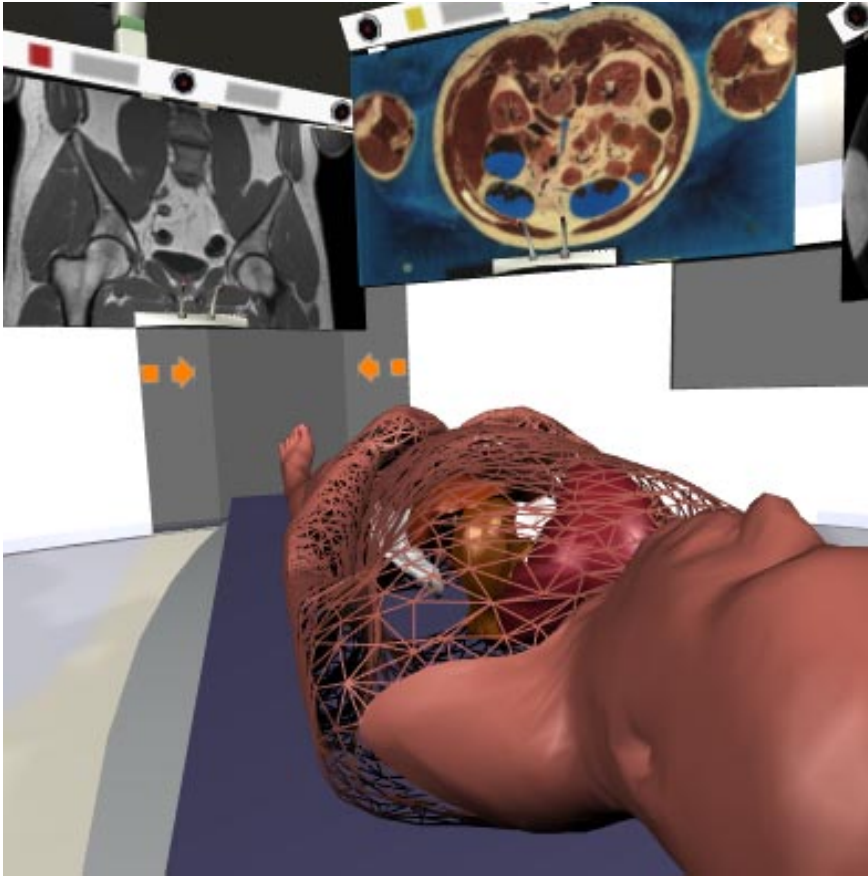


Figure 5: The virtual patient

2.3.1 Shading

In general, *shading* is the realistic display of an object, based on the position, orientation and characteristics of its surface and the light sources illuminating it [26]. The reflective properties of a surface are described with an *illumination model* such as the *Phong* model, which uses a combination of ambient light, diffuse (like paper) and specular (like polished metal) reflections. A key input into these models is the local surface orientation, described by a *normal vector* perpendicular to the surface.

The original marching cubes algorithm calculates the surface normal vectors from the gray level gradients in the data volume [36], described later. Alternatively, the surface normal vectors of the triangles can be used directly. Images produced with these two methods are compared in [73, 93].

2.4 Volume-Based Rendering

In *volume-based rendering*, images are created directly from the volume data. Compared to surface-based methods, the major advantage is that all gray level information which has originally been acquired is kept during the rendering process. As shown by Höhne et al. [38], this makes it an ideal technique for interactive data exploration. Threshold values and other parameters which are not clear from the beginning can be changed interactively. Furthermore, volume-based rendering allows a combined display of different aspects such as opaque and semi-transparent surfaces, cuts, and maximum intensity projections. A current drawback of volume-based techniques is that the large amount of data which has to be handled does not allow real-time

applications on present day computers.

2.4.1 Scanning the Volume

In volume-based rendering, we basically have the choice between two scanning strategies: pixel by pixel (image order) or voxel by voxel (volume order). These strategies correspond to the image and object order rasterization algorithms used in computer graphics [26].

In *image order* scanning, the data volume is sampled on rays along the view direction. This method is commonly known as *ray casting*:

```
FOR each pixel on image plane DO
  FOR each sampling point on associated viewing ray DO
    compute contribution to pixel
```

The principle is illustrated in fig. 6. Along the ray, visibility of surfaces and objects is easily determined. The ray can stop when it meets an opaque surface. Yagel et al. extended this approach to a full *ray tracing* system which follows the viewing rays as they are reflected on various surfaces [109]. Multiple light reflections between specular objects can thus be handled.




Figure 6: Principle of ray casting for volume visualization. In this case, the object surface is found using an intensity threshold.

Image order scanning can be used to render both voxel and polygon data at the same time [54]. Image quality can be adjusted by choosing smaller (oversampling) or wider (undersampling) sampling intervals [53, 70]. Unless stated otherwise, all 3D images shown here were rendered with a ray casting algorithm.

As a drawback, the whole input volume must be available for random access to allow arbitrary view directions. Furthermore, interpolation of the intensities at the sampling points is required. A strategy to reduce computation times is based on the observation that most of the time is spent traversing empty space, far away from the objects to be shown. If the rays are limited to scan the data only within a pre-defined *bounding volume* around these objects, scanning times are greatly reduced [4].

In *volume order* scanning, the input volume is sampled along the lines and columns of the 3D array, projecting a chosen aspect onto the image plane in the direction of view:

```

FOR each sampling point in volume DO
  FOR each pixel projected onto DO
    compute contribution to pixel

```

The volume can either be traversed in back-to-front (BTF) order from the voxel with maximal to the voxel with minimal distance to the image plane, or vice versa in front-to-back (FTB) order. Scanning the input data as they are stored, these techniques are reasonably fast even on computers with small main memories, and especially suitable for parallel processing. So far, ray casting algorithms still offer a higher flexibility in combining different display techniques. However, volume rendering techniques working in volume order are available [105].

2.4.2 Shaded Surfaces

Using one of the described scanning techniques, the visible surface of an object can be determined with a threshold or an object label. For shading, any of the methods developed for the cuberille model, such as distance or distance gradient shading, can be applied.

As shown by Höhne and Bernstein [36], a very realistic and detailed presentation is obtained if the gray level information present in the data is taken into account. Due to the partial volume effect, the gray levels in the 3D neighborhood of a surface voxel represent the relative proportions of different materials inside these voxels. The resulting *gray level gradients* can thus be used to calculate surface inclinations. Unless stated otherwise, this method was used for all 3D images shown here. The simplest variant is to calculate the components of a gradient G for a surface voxel at (i, j, k) from the gray levels g of its six neighbors along the main axes as

$$\begin{aligned}
 G_x &= g(i+1, j, k) - g(i-1, j, k) \\
 G_y &= g(i, j+1, k) - g(i, j-1, k) \\
 G_z &= g(i, j, k+1) - g(i, j, k-1)
 \end{aligned}$$

Scaling G to unit length yields the surface normal [93]. The gray level gradient may also be calculated from all 26 neighbors in a $3 \times 3 \times 3$ neighborhood, weighted according to their distance from the surface voxel [93]. Aliasing patterns are thus almost eliminated.

In the case of very small objects like thin bones, the gray level gradient does not correspond to the actual surface inclination any more. Pommert et al. proposed an *adaptive gray level gradient method* which chooses only 3–6 meaningful neighbors, maximizing the gradient magnitude [73, 93]. This algorithm yields smooth images even for thin objects.

2.4.3 Cut Planes

Once a surface view is available, a very simple and effective method to visualize interior structures is cutting. When the original intensity values are mapped onto the cut plane, they can be better understood in their anatomical context [38]. A special case is selective cutting, where certain objects are excluded (fig. 7).




Figure 7: 3D image of a brain from MRI. Original intensity values are mapped onto the cut planes.

2.4.4 Integral and Maximum Intensity Projection

A different way to look into an object is to integrate the intensity values along the viewing ray. If applied to the whole data volume, this is a step back to the old X-ray projection technique. If applied in a selective way, this *integral projection* is nevertheless helpful in certain cases [38, 93].

For small bright objects such as vessels from magnetic resonance angiography (MRA), *maximum intensity projection* (MIP) is a suitable display technique (fig. 8). Along each ray through the data volume, the maximum gray level is determined and projected onto the image plane [23]. The advantage of this method is that neither segmentation nor shading are needed, which may fail for very small vessels. But there are also some drawbacks: as light reflection is totally ignored, maximum intensity projection does not give a realistic 3D impression. Spatial perception can be improved by rotating the object or by a combined presentation with other surfaces or cut planes [38].

2.4.5 Volume Rendering

Volume rendering is the visualization equivalent to fuzzy segmentation. For medical applications, these methods were first described by Drebin et al. [21] and Levoy [53]. A commonly assumed underlying model is that of a colored, semi-transparent gel with suspended low-albedo (low reflectivity) particles [9]. Illumination rays are partly reflected and change color while traveling through the volume.

Each voxel is assigned a color and an opacity. This opacity is the product of an “object weighting function” and a “gradient weighting function”. The object weighting function is usually dependent of the gray level, but it can also be the result of a more sophisticated fuzzy segmentation algorithm. The gradient weighting function emphasizes surfaces for 3D display. All voxels are shaded, using the gray level gradient method. The shaded values along a viewing ray are weighted and summed up. A somewhat simplified basic equation modeling frontal illumination with a ray casting system is given as follows:




Figure 8: 3D images of head vessels from magnetic resonance angiography (MRA). While the maximum intensity projection (left) shows some more details, spatial perception is much better for thresholding segmentation and gradient-based surface shading (right).

<i>intensity</i>	intensity of reflected light	
<i>p</i>	index of sampling point on ray	(0 ... maximum depth of scene)
<i>l</i>	fraction of incoming light	(0.0 ... 1.0)
α	local opacity	(0.0 ... 1.0)
<i>s</i>	local shading component	

$$intensity(p, l) = \alpha(p) \cdot l \cdot s(p) + (1.0 - \alpha(p)) \cdot intensity(p + 1, (1.0 - \alpha(p)) \cdot l)$$

The total reflected intensity as displayed on a pixel of the 3D image is given as $intensity(0, 1.0)$.

Since binary decisions are avoided in volume rendering, the resulting images are very smooth and show a lot of fine details (fig. 9). Another important advantage is that even coarsely defined objects can be rendered [93].

On the other hand, the more or less transparent images produced with volume rendering methods are often hard to understand so that their clinical use may be limited [93]. Spatial perception can however be improved by rotating the object. Another problem is the large number of parameters which have to be specified to define the weighting functions. Furthermore, volume rendering is comparably slow because weighting and shading operations are performed for many voxels on each ray.

2.5 Transform-Based Rendering

While both surface- and volume-based rendering are operating in a 3D space, 3D images may also be created from other data representations. One such method is *frequency domain rendering*, which creates 3D images in Fourier space, based on the projection-slice theorem [96]. This method is very fast, but the resulting images are limited to rather simple integral projections (see above).



Figure 9: Volume rendered image of a child with a craniosynostosis (trigonocephalon) from CT. Semi-transparent visualization shows a lot of different aspects such as skin and bone surface, but spatial perception is quite difficult.

A more promising approach are *wavelet transforms*. These methods provide a multi-scale representation of 3D objects, with the size of represented detail locally adjustable. Amount of data and rendering times may thus dramatically be reduced. Application to volume visualization is shown in [63, 64].

2.6 Multimodality Matching

For many clinical applications, it is desirable to combine information from different imaging modalities. For example, for the interpretation of PET images which show only physiological aspects, it is important to know the patient's morphology, as shown in MRI. In general, different data sets do not match geometrically. It is therefore required to transform one volume with respect to the other. This process is also known as *registration*.

The transformation may be defined using corresponding *landmarks* in both data sets [97]. In a simple case, external markers attached to the patient are available which are visible on different modalities. Otherwise, arbitrary pairs of matching points may be defined. A more robust approach is to interactively match larger features such as surfaces, or selected internal features such as the AC-PC line (anterior / posterior commissure) in brain imaging [84]. All these techniques may also be applied in scale-space at different levels of resolution [62].

In a fundamentally different approach, the results of a registration step are evaluated at every point of the combined volume using *voxel similarity measures*, based on intensity values [91, 104]. Starting from a coarse match, registration is achieved by adjusting position and orientation until the mutual information between both data sets is maximized. Since these methods are fully automatic and do not rely on a possibly erroneous definition of landmarks, they are increasingly considered superior.

2.7 Intervention Rehearsal

So far, we have focused on merely visualizing the data. A special case is to move the camera inside the patient for simulated endoscopy [27]. Besides in education, potential applications are in non-invasive procedures, such as gastrointestinal diagnosis.

A step further is to manipulate the data at the computer screen for surgery simulation. These techniques are most advanced for craniofacial surgery where a skull is dissected into small pieces, and then rearranged to achieve a desirable shape. Several systems have been designed which allow the user to interactively draw closed curves onto the screen, which are interpreted as cuts into the volume [66, 110]. The resulting segments can be inspected from other view directions, and individually moved and rearranged in 3D space.

While these systems are based on binary data, Pflesser et al. developed an algorithm which handles full gray level volumes [69]. Thus, all features of volume-based rendering, including cuts and semi-transparent rendering of objects obscuring or penetrating each other, are available.

2.8 Image Quality

For applications in the medical field, it is mandatory to assure that the 3D images show the true anatomical situation, or at least to know about their limitations.

A common approach for investigating *image fidelity* is to compare 3D images rendered by means of different algorithms. This method, however, is of limited value since the truth usually is not known. A more suitable approach is to apply volume visualization techniques to simulated data [57, 73, 93], and to data acquired from corpses [22, 32, 67, 71, 76]. In both cases, the actual situation is available for comparison. Using the first technique, the accuracy e.g. of different shading algorithms has been shown. Results of the latter studies include visibility of sutures or fracture gaps as a function of acquisition parameters and object size.

Another aspect of image quality is *image utility*, which describes whether an image is really useful for a viewer with respect to a certain task. Investigations of 3D image utility in craniofacial surgery may be found in [2, 89, 98].

Part II

Applications to Medical Data

3 Radiological Data

At first glance, one might expect diagnostic radiology to be the major field of application for volume visualization. This is however not always the case. One of the reasons is clearly that radiologists are especially skilled in reading cross-sectional images. Another reason is that many diagnostic tasks such as tumor detection and classification can well be done from tomographic images. Furthermore, 3D visualization of these objects from MRI requires robust segmentation algorithms which are not yet available. In [78] one can find a selection of advanced applications and "success stories" from the application of methods reported here.

The situation is generally different in all fields where therapeutical decisions have to be made by non-radiologists on the basis of radiological images [37, 111]. A major field of application for volume visualization methods is *craniofacial surgery* [2, 19, 55, 112]. Volume visualization not only facilitates understanding of pathological situations, but is also a helpful tool for planning optimal surgical access and cosmetic results of an intervention. A typical case is shown in fig. 9. Dedicated procedures for specific disorders have been

developed, which are now in routine application [89].

Another important field of application is *traumatology* (fig. 10). Due to the emergency situation, planning times are usually very short. With new faster imaging modalities available and computing power ever increasing, volume visualization techniques are being introduced for difficult cases [25].



Figure 10: 3D image of a pelvis with multiple fractures from CT in a semi-frontal view.

An application that is becoming more and more attractive with the increasing resolution and specificity of MRI is *neurosurgery planning*. Here the problem is to choose a proper access path to a lesion. 3D visualization of brain tissue from MRI and blood vessels from MRA before surgical intervention allows the surgeon to find a path with minimal risk in advance [18, 70]. In combination with a 3D coordinate digitizer, the acquired information can be used to guide the surgeon during the intervention [1]. In conjunction with functional information from PET images, localization of a lesion is facilitated. The state of the art in *computer-integrated surgery* is presented in [92].

Another important application that reduces the risk of a therapeutical intervention is *radiotherapy planning*. Here, the objective is to focus the radiation as closely as possible to the target volume, while avoiding side effects in healthy tissue and radiosensitive organs at risk. 3D visualization of target volume, organs at risk and simulated radiation dose allows an iterative optimization of treatment plans [41, 46, 83, 86].

Applications apart from clinical work include *medical research and education*. In the current Decade of Brain Research, exploring and mapping brain functions is a major issue. Volume visualization methods provide a framework to integrate information obtained from such diverse sources as dissection, functional MRI, or magnetoencephalography [95].

The knowledge for the interpretation of the 3D pictures described so far has still to come from the viewer. In contrast, the 3D brain atlas VOXEL-MAN/brain shown in figs. 7 and 12 is based on an *intelligent volume*, which has been prepared from an MRI data set [40, 72, 88]. It contains spatial and semantic descriptions of morphology, function, and blood supply. The brain may be explored on the computer screen in a style close to a real dissection, and queried at any point. Other atlases are currently in preparation (fig. 3). Beyond




Figure 11: Gradual transition between surface and MIP visualization of a gamma camera dataset of the pelvis. Heart, kidneys, liver and spleen are visible. Three haemangiomas can be seen in the MIP mode

education, such atlases are also a powerful aid for the interpretation of clinical images [84].

If high resolution cryosections such as those created in the *Visible Human Project* of the National Library of Medicine [90] are used, even more detailed and realistic atlases can be prepared [85, 94]. An example is shown in fig. 13.

4 3D Ultrasound

4.1 Introduction

3D ultrasound is a very new and most interesting application in the area of 'tomographic' medical imaging, able to become a fast, non-radiative, non-invasive, and inexpensive volumetric data acquisition technique with unique advantages for the localisation of vessels and tumours in soft tissue (spleen, kidneys, liver, breast etc.). In general, tomographic techniques (CT, MR, PET etc.) allow for a high anatomical clarity when inspecting the interior of the human body. In addition, they enable a 3D reconstruction and examination of regions of interest, offering obvious benefits (reviewing from any desired angle, isolation of crucial locations, visualization of internal structures, 'fly-by', accurate measurements of distances, angles, volumes etc.).

The physical principle of ultrasound is as following [61]: sound waves of high frequency (1–15 MHz) emanate from a row of sources that are located on the surface of a transducer which is in direct contact with the skin. The sound waves penetrate the human tissue travelling with a speed of 1450–1580 m/s, depending upon the type of tissue. The sound waves are reflected partially if they hit an interface between two different types of tissue (e.g. muscle and bone). The reflected wavefronts are detected by sensors (microphones) located next to the sources on the transducer. The intensity of reflected energy is proportional to the sound impedance difference of the two corresponding types of tissue and depends on the difference of the sound impedances Z_1

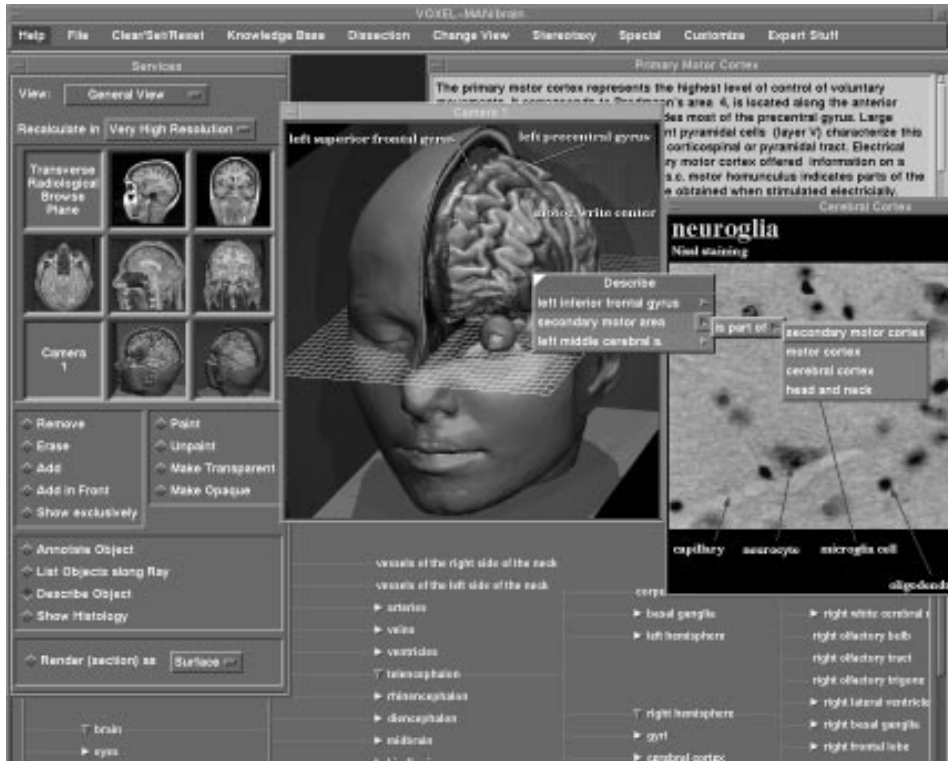


Figure 12: Anatomy teaching by dissection at the computer screen: user interface of VOXEL-MAN/brain. The volume may be arbitrarily rotated and cut. By pointing to a visible surface, the selected region can be highlighted and annotated. Objects to be displayed may also be selected from a knowledge base.

and Z_2 :

$$I_r = I_e \cdot \frac{\left(1 - \frac{Z_2}{Z_1}\right)}{\left(1 + \frac{Z_2}{Z_1}\right)} \quad (1)$$

An image of the interior structure can be reconstructed based upon the total travelling time, the (average) speed, and the energy intensity of the reflected waves. The resulting 3D images essentially represent hidden internal “surfaces”. The principle is similar to radar with the difference being that it uses mechanical instead of electromagnetic waves.

4.2 Collecting 3D-Ultrasound Data

In contrast to the common 2D case where a single image slice is acquired, 3D ultrasonic techniques cover a volume within the body with a series of subsequent image slices. The easiest way to collect 3D-ultrasound data is to employ a Kretz Voluson 530 device. This is a commercially available device allowing for a direct acquisition of a hole volume area instead of a single slice. The principle of the Kretz device is based on a mechanical movement of the transducer during acquisition along a rotational or sweep path (see figure 15)

The advantage of the Kretz system lies in its high precision and commercial availability. Its disadvantage is that the rather high system price makes it somehow difficult to purchase for physicians. The alternative is a free-hand scanning system allowing the upgrade of virtually any existing conventional (2D) ultrasound system to full 3D-capabilities. Such an update can be done by exclusively external components and hence

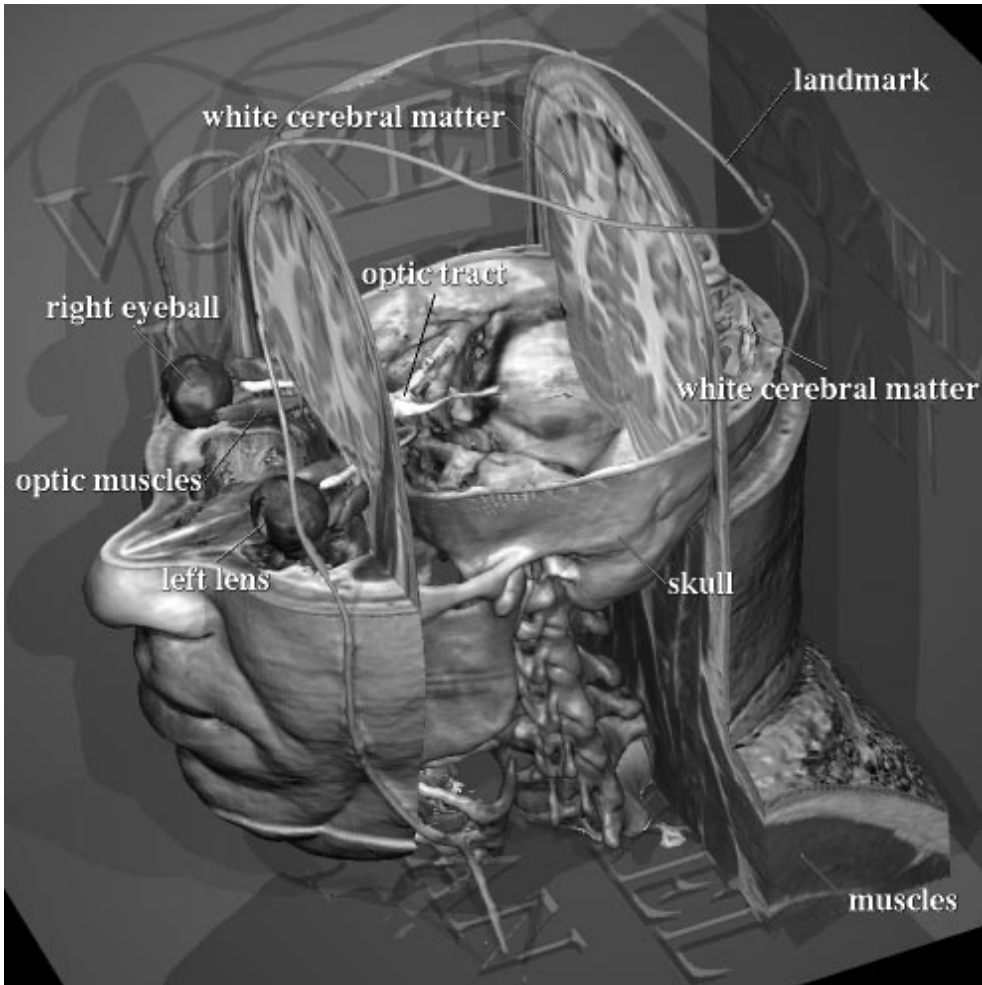


Figure 13: Dissection of the Visible Human. Used in a state of the art visualization environment, this data represents a new quality of anatomical imaging.

does not require any manipulation of the existing hardware and software configuration. After the upgrade the ultrasound equipment can be operated in both, 2D as well as 3D mode almost simultaneously. Switching from the 2D to the 3D mode requires only a mouse-click. As a result, the familiar 2D examination procedure remains unchanged, and the physician can switch on the 3D mode only there when this is necessary. The system architecture is illustrated in the figure below. The upgrade requires the employment of two external components:

1. A 6-degrees-of-freedom (6DOF) tracking system for the transducer. Such a tracking system is mounted on the transducer and follows very precisely its position and orientation in 3D space. Thus, each 2D image is associated with a corresponding position and orientation co-ordinates. The physician can now move the transducer free-hand over the region under examination. In the market there exist several different types of 6DOF tracking systems: mechanical arms, electromagnetic trackers, camera-based trackers (infrared or visible light).
2. A image digitalisation and volume rendering system. This component consist from a frame grabber, a workstation or PC with sufficient memory and processor power, a serial interface and the usual periph-




Figure 14: The principal function of ultrasound


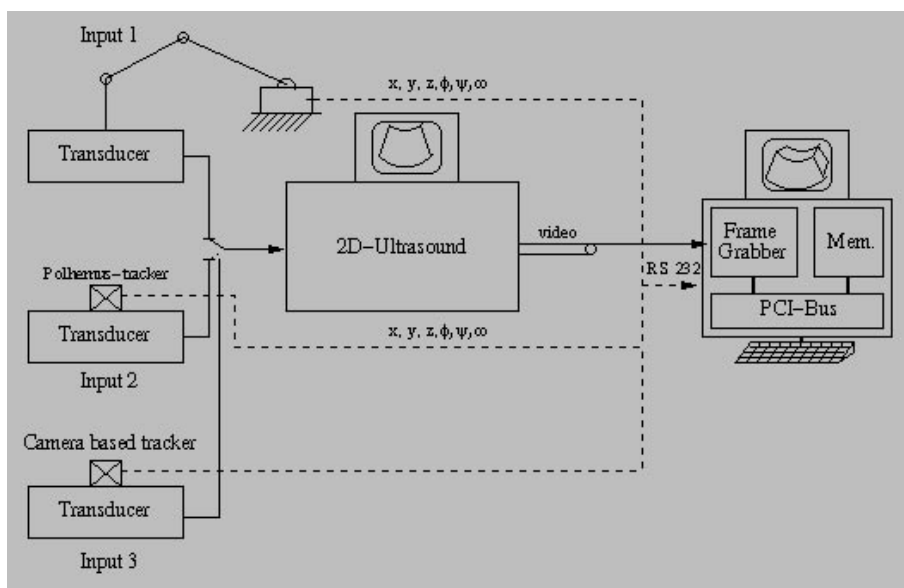


Figure 15: Different mechanical scanning methods

eral devices (monitor, mouse, printer etc.). The video output of the 2D ultrasound machine is connected to the frame grabber and the 6DOF tracker to the serial input. Every 2D image presented on the ultrasound screen is digitised in real-time and stored together with its corresponding tracker co-ordinates in the memory. After finishing the scanning procedure, all acquired 2D slices are combined into a 3D volume sample of the examined area. This volume dataset is then further processed.



4.3 Visualization of 3D Ultrasound

One of the major reasons for the limited acceptance of 3D ultrasound to date is the complete lack of an appropriate visualisation technique, able to display clear surfaces out of the acquired data. The very first approach was to use well known techniques, used for MRI and CT data to extract surfaces. Such techniques, reported in more detail in the first part of this Report, include binarization, iso-surfacing, contour connecting, marching cubes, and volume rendering either as semi-transparent cloud, or as fuzzy gradient shading [52]. Manual contouring is too slow and impractical for real-life applications. Unfortunately, ultrasound images possess several features causing all these techniques to fail totally. The general appearance of a volume rendered 3D ultrasound dataset is that of a solid block covered with 'noise snow' (fig. 16 right). The most important of these features as reported in [79] and [81] are:

1. significant amount of noise and speckle
2. much lower dynamic range as compared to CT or MR
3. high variations in the intensity of neighbouring voxels, even within homogeneous tissue areas
4. boundaries with varying grey level caused by the variation of surface curvature and orientation to the sound source
5. partially or completely shadowed surfaces from objects closer and within the direction of the sound source (e.g. a hand shadows the face)
6. the regions representing boundaries are not sharp but show a width of several pixels
7. poor alignment between subsequent images (parallel-scan devices only)
8. pixels representing varying geometric resolutions depending on the distance from the sound source (fan-scanning devices only)

The next idea in dealing with ultrasound data was to improve the quality of the data during a pre-processing step, i.e. prior to reconstruction, segmentation and volume rendering. When filtering medical images, a trade-off between image quality and information loss must always be taken into account. Several different filters have been tested: 3D-Gaussian for noise reduction, 2D-speckle removal for contour smoothing, 3D-median for both noise reduction and closing of small gaps caused by differences in the average luminosity between subsequent images [79]; other filters such as mathematical topology and extended threshold-based segmentation have been tested as well. The best results have been achieved by combining Gaussian and median filters (see image 17).




Figure 16: Left a grey image of the liver, middle the corresponding opacity values, right a volume rendered dataset. Note the high opacity values along the interface between data and empty space (middle) causing a solid 'curtain' obscuring the volume interior (right)




Figure 17: Volume rendering after off-line 3D-median and 3D-Gaussian filtering. Form left to right: unfiltered and median with a width of 3^3 , 5^3 and 7^3 . In the lower row the same data after additional Gaussian filtering with a width of 3^3

However, pre-processing of large datasets (a typical 3D-volume has a resolution of 256^3 voxels) requires several minutes of computing, reduces the flexibility to interactively adjust visualization parameters, and aliases the original data. For solving these problems interactive filtering techniques based on multi-resolution analysis and feature extraction have been developed, allowing a user-adjustable, on-line filtering within a few seconds and provide an image quality comparable to the off-line methods ([81], see image 18).

In order to remove artifacts remaining in the image after filtering, semi-automatic segmentation has been applied because of the general lack of a reliable automatic technique. A segmentation can be provided by using the mouse to draw a few crude contours (see [79] for more details).

The diagnostic value of surface reconstruction in prenatal diagnosis so far has to be seen in the routine detection of small irregularities of the fetal surface, such as cheilognathopalatoschisis or small (covered) vertebral




Figure 18: On-line filtering of the face of a fetus. This filtering is completed within less than 5 seconds




Figure 19: Fetal face before (left) and after (middle) removing the right hand and the remaining artefacts (right)

defects as well as in a better spatial impression of the fetus as compared to the 2D-imaging. A useful side-effect is a psychological one, as the pregnant woman gets a more plastic impression of the unborn ([35], [5]).




Figure 20: On-line mixing between surface and MIP models. This operation is performed in real-time

Image 21 compares an image reconstructed from data acquired in the 25th week of pregnancy with a photo of the baby 24 hours after birth. The resolution of the data was $256 \times 256 \times 128$ (8 Mbytes); the time for volume rendering one image with a resolution of 300^2 pixels is about 1 second on a Pentium Pro 200 doppelprocessor PC.




Figure 21: Comparison of a volume reconstructed from 3D-ultrasound data acquired during the 25th pregnancy week (3 1/2 months before birth) with a photograph of the same baby taken 24 hours after birth

Image 22 shows several other examples of fetal faces acquired in the Mannheim Clinic. It is important to note that these datasets have been acquired under routine clinical conditions and therefore they can be regarded as to be representative. In average 80% of the acquired volumes can be reconstructed within ca. 10 minutes with an image quality comparable to that shown here. All cases where the fetus was facing the abdominal wall could be reconstructed successfully.




Figure 22: Six different examples of fetal faces acquired under daily clinical routine conditions

Under clinical aspects further work should be aimed towards a better distinction and automatic separation of surfaces lying close together and showing relatively small grey-scale differences. The reconstruction of

surfaces within the fetus, e.g. organs, is highly desired. Surface properties of organs, but also of pathological structures (ovarian tumors etc.) might give further information for the assessment of the dignity of tumors.

5 3D Cardiac Reconstruction from 2D Projections

Different imaging modalities are applied in order to acquire medical data. In terms of the human heart 3D tomographic imaging techniques are not yet suitable for resolving neither moving coronary arteries nor the changing volume of the heart-ventricles.

The golden standard for diagnosis of coronary artery disease or volumetry is x-ray angiography, recently combined with intra-vascular ultrasound (IVUS) [51]. The main benefit of this technique is the high spatial and temporal resolution as well as high image contrast.

For treatment planning of angioplasty or bypass surgery or for volumetry sequences of x-ray images are traditionally acquired and evaluated. Despite the high quality of angiograms an exact judgment of pathological changes (eg. stenosis) requires a large amount of experience on the part of the cardiologist.

In order to improve the diagnostic accuracy, 3D reconstruction from 2D coronary angiograms appears desirable [101]. In general two different approaches can be distinguished. The stereoscopic or multiscopic determination of ray intersections is a method which makes it necessary to identify correspondent features within different images. If this correspondence is not possible to be established, back-projection techniques [24] are more suitable.

The choice of using either the stereoscopic or the back-projection approach mainly depends on the following criteria:

Number of images For stereoscopic approach at least two images are necessary to perform the reconstruction. In order to achieve good results by using back-projection techniques more than twenty images are necessary.

Relative orientation A small relative orientation results in a low accuracy for both stereoscopic and back-projection techniques. Nevertheless the necessity of a large parallax angle is higher for back-projection techniques.

Morphology In order to reconstruct objects which are composed by a number of small structured parts stereoscopic techniques are more appropriated. On the other hand large objects with low structure are easier to reconstruct by back-projection techniques.

Occluding objects Occluding objects cause problems when using stereoscopic methods. In contrast, back-projection techniques are able to separate different objects which are laying on the same projection ray.

Since the choice of the right technique strongly depends on the current application both approaches will be described briefly within the following sections.

5.1 Reconstruction of Coronary Vessels

In this section a method of reconstructing the 3D appearance of the coronary arteries, based on a sequence of angiograms, acquired by rotating a mono-plane system around the heart will be described. In order to determine the exact phase of the heart cycle for each image an ECG is recorded simultaneously. In order to minimize user interaction and a priori knowledge introduced into the reconstruction process [30] a new method has been developed and implemented. The technique requires a minimum of user interaction limited

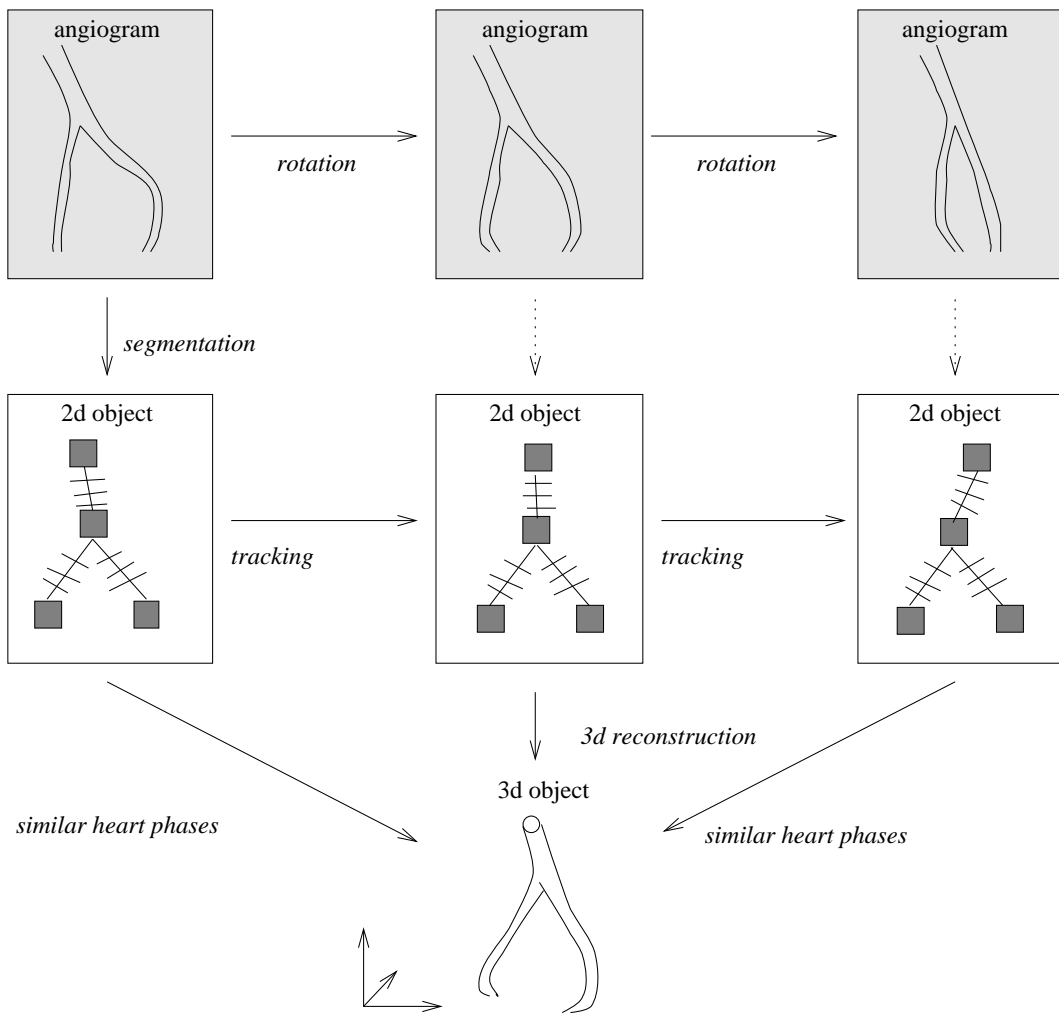


Figure 23: Processing steps used to reconstruct the 3D geometry of the coronary vessels

to the segmentation of vessels in the initial image of each angiographic sequence. The segmentation result is exploited in the entire series of angiograms to track each individual vessel.

In contrast to the assumption for 3D reconstruction of objects from multiple projections, coronary arteries are not rigid. Due to the deterministic nature of the mobility of the heart with respect to the phase of the heart motion, distinct images are used, showing the heart at the same phase of the cardiac cycle. The different processing steps, used for reconstructing the 3D geometry of the vessel are shown in figure 23 and discussed below [34].

In order to separate the vessel tree to be reconstructed, the image has to be segmented. The major drawback of most of the existing segmentation algorithms are either a very limited amount of variation in the input data amenable to processing by a fully automatic algorithm, and the necessity of an extensive user-assistance.

The approach leads to a compromise in which the user only identifies a very small number of points interactively. The segmentation process is separated into the detection of the vessel centerline and evaluation of the vessel contour. The algorithm works with a cost-minimizing A^* search tree [100], [65], which proved to be robust against noise and may be fully controlled by the user. The obtained structure is tracked over the angiographic sequence by Snakes.

Reconstruction is based on the extracted vessel tree structures, the known relative orientation (i.e. the angle)

of the projections, and the imaging parameters of the x-ray system. The 3D reconstruction is performed from images of identical heart phases. It begins with the two projections of the same phase, defining the largest angle. The obtained result is afterwards improved by introducing additional views. Applying a 3D optimization techniques the shape of a 3D Snake is adapted according to multiple 2D projections [33].

The obtained 3D structure can be either visualized by performing a volume rendering, or in order to be presented within a VR-systems it can be transferred into a polygonal representation.

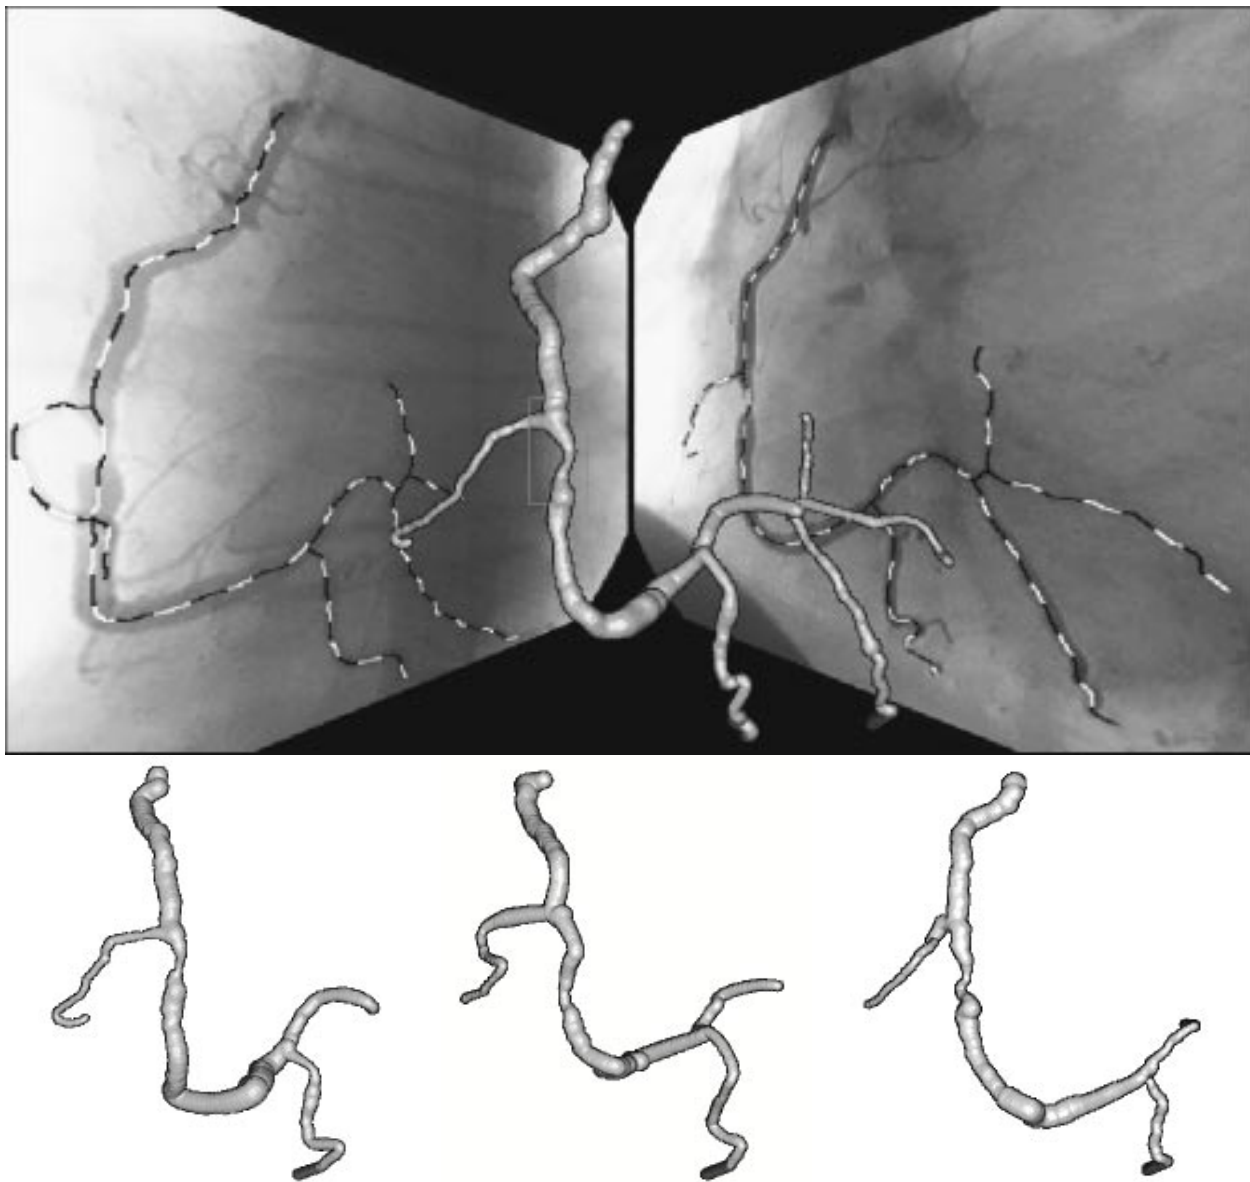


Figure 24: Reconstructed vessels rendered by InViVo (*top*: combined presentation of the volume rendered reconstruction result and angiograms *bottom*: some frames of the 3D movement simulation)

Besides the 3D geometry of the coronary vessels, the trajectories of distinct points of the vessels are determined during the tracking process. As a result these trajectories can be used to simulate the movement of the vessel, caused during the heart beat (Figure 24 bottom row).

5.2 Reconstruction of Ventricles

Beside the stereoscopic or multiscopic feature based approach, the 3D structure can also be obtained using densitometric information. This technique, also known as back-projection method, does not need any a priori knowledge or image segmentation.




Figure 25: Angiograms acquired by the bi-plane x-ray system

Similar to CT, the 3D information is obtained by determining the intensity of a volume element according to the density of the imaged structure. The intensity of each pixel within the angiogram correlates to the amount of x-ray energy which is received at the image amplifier. This energy depends on the density and the absorption capabilities of the traversed material. As a result a pixel represents the sum of the transmission coefficients of the different materials which are pierced by the x-ray. For homogeneous material and parallel monochromatic x-rays the image intensity can be described by the rule of *Lambert-Beer* [8]:

$$I = I_0 e^{-\mu v d} \quad (2)$$

- I : Image intensity
- I_0 : Initial intensity
- μ : Absorption-coefficient of the structure
- v : Density of the structure
- d : Thickness of the structure

If the x-ray travels through a material with varying densities, this equation has to be split into parts with constant density. The total amount of transmitted intensity is the sum of these different parts.

$$I = I_0 e^{-\sum_i \mu_i v_i d_i} \quad (3)$$

To improve the image quality contrast agent is injected during the acquisition process. For this purpose a catheter is positioned in front of the ventricles (see figure 25).

Applying the back-projection technique the distribution of the coefficients can be determined. During the acquisition process the x-ray system is rotated around the center of the heart (see figure 26).

In order to reconstruct the appropriate intensities of the heart, all the images are translated into the center of rotation (see figure 27), thus according to the amount of images a cylinder is defined by a number of sampling




Figure 26: Acquisition of different angiograms by rotating a bi-plane x-ray system around the center of the heart

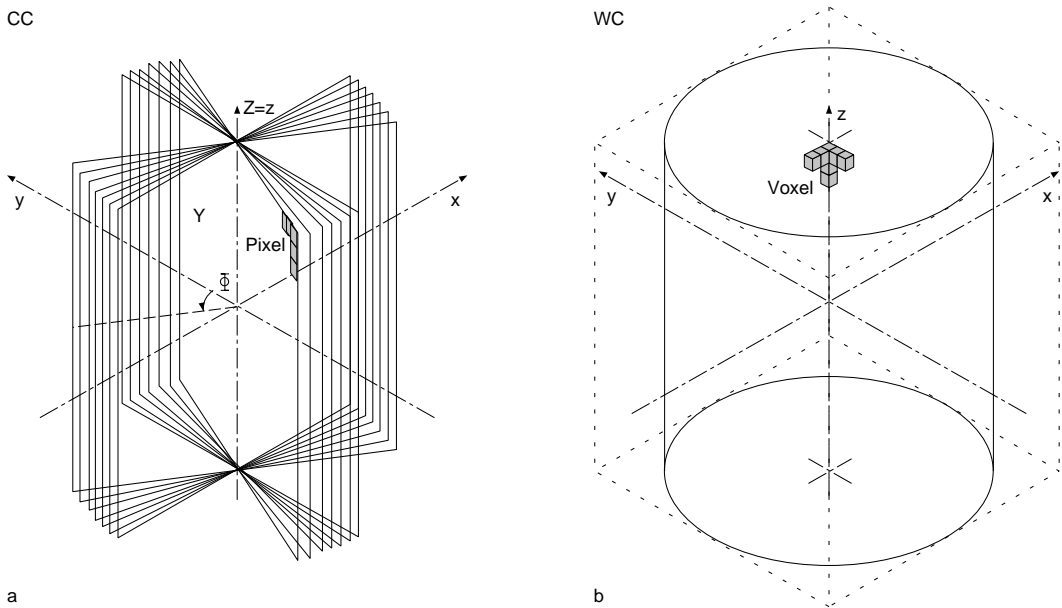


Figure 27: Translation of the angiograms in order to determine the voxel intensities

planes. The complete volume of the cylinder can now be determined. Therefore all the rays starting from the x-ray source and intersecting a distinct voxel are accumulated and weighted according to the intensity of the different planes. Continuing this process for all the voxels of the cylinder, by taking the projection geometry into account by introducing a cone-filter [42], the intensity of each cylinder voxel can be determined. The obtained volume data can be visualized using a volume rendering technique and segmented by Snakes.

6 Visualization of Laser Confocal Microscopy Datasets

Structures in the microscopic scale nerve cells, tissue and muscles, blood vessels etc. show beautiful, complex, and mostly still unexplored patterns usually with higher complexity than those of organs. In order to understand the spatial relationship and internal structure of such microscopic probes, tomographic series of slices are required in analogy to the tomographies used for organs and other macroscopic structures.

Laser confocal microscopy is a relatively new method allowing for a true tomographic inspection of micro-




Figure 28: Volume rendering of the intensities obtained by the back-projection technique

scopic probes. The method operates according to a simple, basic principle [20].




Figure 29: Principle of Laser Confocal Microscopy

Visible or ultra violet laser emission is focused on the first confocal pinhole and then onto the specimen as a diffraction-limited light spot, see fig. 29. The primary incident light is then reflected from particular voxel elements or emitted from fluorescent molecules excited within it. Emissions from the object return along the primary laser light pathway and depart from it by lateral reflection from (or passage through, depending on the instrument) a dichroic mirror onto the second confocal pinhole. This aperture is confocal with the in-focus voxel elements in the specimen. The virtual elimination by defocusing of all distal and proximal flanking emissions at this physical point assures that the light passing onto the detector, a sensitive photodetector or

camera, is specifically derived from in-focus object voxels with a resolution, e.g. in the Leica instrument, approaching 200 to 400 nm in the x/y and z directions, respectively. In order to image the entire object, the light spot is scanned by a second mirror in the x/y plane in successive z sections by means of a precision stage motor. Rapid scanning preserves fluorescent intensity but must be reconciled with image quality. The storage, retrieval and manipulation of light intensity information from the object makes static and dynamic 3-D imaging possible.

Although not perfect, the new method shows several significant benefits as compared to the traditional procedures. The most important of them are: true tomographic method, significant freedom in choosing slice thickness and size, trivial registration of slices, very fast and easy in operation, capable of acquiring in-vivo cells as well as static or dynamic structures, non-destructive. Lastly, by using different types of laser and fluorophore materials, different spatially overlapping structures can be visualized and superimposed within the same probe.

The data acquired with Laser Confocal Microscopy (LCM) show several characteristics requiring specialized treatment in order to make the method applicable:

1. Large data size. Typical datasets have a resolution of $512^2 \times 64$ pixels. These pixels are colored, thus a typical RGB dataset requires some 50 Mbytes of memory. Obviously, datasets of this size require efficient processing methods.
2. Low contrast, low intensity gradients, bad signal to noise ratio. These characteristics make a straightforward segmentation between the structures of interest and the background (e.g. by using thresholding, region growing, homogeneity, color differences etc.) impossible. All the methods listed above apply more or less binary decision criteria if a pixel/voxel belongs to the structure or not. Such criteria typically fail when used with signals showing the characteristics listed above.
3. Unequal resolutions in the plane and the depth directions. Thus, a visualization method has to be able to perform with "blocks" or unequal size lengths instead with cubic voxels. Re-sampling of the raw data to a regular cubic field will further reduce the signal quality, introduce interpolation artifacts and generate an even larger dataset, probably too large to be handled with conventional computers.
4. Regarding the quality, artifacts have to be avoided as far as possible. Introducing artifacts in an unknown structure will often have fatal effects on their interpretation, since the human observer does not always have the experience for judging the correctness or the fidelity of the presented structures. As an example, an obvious artifact caused by bad parameter settings of the software during the visualization of human anatomy (e.g. of a head) is immediately detected by the observer, since the human anatomy is well known and such artifacts are trivially detected. This is not the case when inspecting an unknown dataset.
5. Choosing the "correct" illumination model (e.g. MIP, semi-transparent, surface etc.) has a significant impact on the clarity and information content of the visualization. Again, due to the lack of experience such a decision is typically much more difficult than in the case of anatomic tomographic data.
6. The speed of visualization becomes the most crucial issue. The visualization parameters have to be adjusted in an interactive, trial-and-error procedure. This can take a very long time if, e.g., after an adjustment the user has to wait for several minutes to see the new result. Furthermore, inspection of new, unknown structures require rapid changing of directions, illumination conditions, visualization models, etc.. Looping and stereo images are of enormous importance for understanding unknown, complicated spatial structures.

The main requirement here is to employ a fast volumetric method allowing interactive and intuitive parameter settings during the visualization session. Detailed results of the employed volume visualization are reported

in [80]. Figures 31 and 30 present a microscopic preparation of the tubular structure of a cat retina. The dataset consist of $335 \times 306 \times 67$ voxels, each with a dimension of $0.16^2 \times 0.2\mu m$. The first image presents the extra-cellular component of the blood vessel. The vessel diameter before the branch point is $19 \mu m$. The second image shows the wire-like structure of the astrocyte cytoskeleton. Both datasets originate from the same probe. In all subsequent images the difference of the visualization between slicing, MIP, surface and semi-transparent methods is shown.




Figure 30: The wire-like structure of the astrocyte cytoskeleton of the same probe. Resolution of $335 \times 306 \times 67$ voxels with a size of $0.16^2 \times 0.2\mu m$. Upper left a single slice, right surface reconstruction; lower left MIP, right transmission illumination models

The next dataset 32 shows the complicated structure of nerve cells networks. The resolution of the dataset is with 25 Mbytes ($512^2 \times 100$ voxels) large. As one can see on image 32 upper left, single slices are not able to provide full understanding of the complicated topology. The three other images show in much better detail the internal structure of the cell network.

Laser confocal microscopy plays a fundamental role for gathering in-vivo data about not only static, but also dynamic structures, i.e. structures existing typically only within living cells and for a very short period of time (e.g. for a few seconds). Such structures are common in several biological applications. In the case referred here we present temporary structures formed by polymerized actin, a structure necessary for cell movements.

Figures 33 and 34 demonstrate the importance of LCM data visualization for detecting unknown structures. In this case, we studied actin filaments in *Dictyostelium* amoebae with time periods ranging from 10 to 100 seconds. The data resolution is $512 \times 484 \times 43$ voxels = 10 Mbytes. Note the structure of the surface visible in the "surface volume rendering" image. These structures are hardly visible and therefore difficult to detect when regarding individual slices.

7 Conclusions

Medical volume visualization has come a long way from the first experiments to the current, highly detailed renderings. As the rendering algorithms are improved and the fidelity of the resulting images is investigated,




Figure 31: The extra-cellular component of a retina blood vessel of a cat. Resolution of $335 \times 306 \times 67$ voxels with a size of $0.16^2 \times 0.2\mu m$. Upper left a single slice, right surface reconstruction; lower left MIP, right transmission illumination models




Figure 32: The complicated topology of nerve cells networks. Resolution $512^2 \times 100$ voxels = 25 MBytes. Upper left a single slice, right surface reconstruction; lower left MIP, right transmission illumination models

3D images are not just pretty pictures, but a powerful source of information for research, education, and patient care. In certain areas such as craniofacial surgery or traumatology, volume visualization is increasingly becoming part of the standard preoperative procedures. New applications such as 3D cardiology, 3D Ultrasound and laser confocal microscopy are becoming more and more popular. Further rapid development of




Figure 33: F-actin structures in *Dictyostelium amoebae*, Resolution $512 \times 484 \times 43$ voxels = 10 Mbytes. Upper left a single slice, right surface reconstruction; lower left MIP, right transmission illumination models




Figure 34: F-actin structures. Resolution $512 \times 484 \times 70$ voxels = 16.5 Mbytes. Upper left a single slice, right surface reconstruction; lower left MIP, right transmission illumination models

volume visualization methods is widely expected [45].

A number of problems still hinder an even broader use of volume visualization in medicine. First, and most importantly, the segmentation problem is still unsolved. It is no coincidence that volume visualization is most accepted in all areas where clinicians are interested in bone from CT. Especially for MRI, however, automatic

segmentation methods are still far from being generally applicable, while interactive procedures are much too time-consuming. As has been shown, there is research in different directions going on; in many cases, methods have already proven valuable for specific applications.

The second major problem is the design of a user interface which is suitable in a clinical environment. Currently, there is still a large number of rather technical parameters for controlling segmentation, matching, shading, and so on. Acceptance in the medical community will certainly depend heavily on progress in this field.

Third, current workstations are not yet able to deliver 3D images fast enough. For the future, it is certainly desirable to interact with the workstation in real time, instead of just looking at static images or pre-calculated movies. However, with computing power further increasing, this problem will be overcome in just a few years even on low-cost platforms.

As has been shown, a number of applications based on volume visualization are becoming operational, such as surgical simulation systems and three-dimensional atlases. Another intriguing idea is to combine volume visualization with *virtual reality* systems, which enable the clinician to walk around or fly through a virtual patient, see fig. 35 [77], [49]. In *augmented reality*, images from real and virtual world are merged to guide the surgeon during an intervention [6].




Figure 35: OP 2000: The Operation Theatre of the Future

Integration of volume visualization with virtual reality and robotics towards *computer-integrated surgery* will certainly be a major topic in the coming decade [92, 103, 77].

Acknowledgements

Georgios Sakas wants to thank Axel Hildebrand, Stefan Großkopf, Jürgen Jäger, Rainer Malkewitz and Stefan Walter who did a lot of work within the different sections described above, and Peter Plath and Mike Vicker from the University of Bremen for the laser confocal microscopy data. Furtherly, data have been provided by Kretztechnik, Visible Human Project, Deutsche Klinik fuer EDiagnostik Wiesbaden, H.T.M. van




Figure 36: Virtual Arthroscopy: using VR and force feed-back for training surgical procedures

der Voort, Department of Molecular Biology, University of Amsterdam, Herbert Stüttler and Martin Hoppe, Leica Lasertechnik GmbH, State Hospital of Darmstadt, Rolf Ziegler (fig. 36).

Andreas Pommert is grateful to Prof. Karl Heinz Höhne, to his colleagues Bernhard Pflesser, Martin Riemer, Thomas Schiemann, Rainer Schubert and Ulf Tiede for providing many of the examples and images presented here, and to the students Jan Freudenberg, Sang-II Kim, Henning Krämer, Kay Priesmeyer, Christian Seebode and Frank Wilmer for many discussions and practical assistance. Tomographic raw data were kindly provided by Siemens Medical Systems, Erlangen (figs. 3, 4, 7, 8, 12), and the National Library of Medicine (fig. 13). Applications are in cooperation with Wolf-Joachim Hölzje, Dept. of Craniofacial Surgery (fig. 9), Werner Lierse, late head of the Dept. of Anatomy (figs. 7, 12), Ernst Richter, Dept. of Pediatrics (fig. 3), and Jürgen Wening, Allgemeines Krankenhaus Altona (fig. 10).

References

- [1] L. Adams, W. Krybus, Dietrich Meyer-Ebrecht, R. Rueger, J. M. Gilsbach, Ralph Moesges, and G. Schloendorff. Computer-assisted surgery. *IEEE Comput. Graphics Appl.*, 10(3):43–51, 1990.
- [2] M. E. Alder, S. T. Deahl, and S. R. Matteson. Clinical usefulness of two-dimensional reformatted and three-dimensionally rendered computerized images: Literature review and a survey of surgeons' options. *J. Oral Maxillofac. Surg.*, 53(4):375–386, 1995.
- [3] L. K. Atata, A. P. Dhawan, J. P. Broderick, M. F. Gaskil-Shipley, A. V. Levy, and N. D. Volkow. Three-dimensional anatomical model-based segmentation of MR brain images through principal axes registration. *IEEE Trans. Biomed. Eng.*, 42(11):1069–1078, 1995.
- [4] R. S. Avila, L. M. Sobierajski, and Arie E. Kaufman. Towards a comprehensive volume visualization system. In *Proc. Visualization '92*, pages 13–20, Boston, MA, 1992.
- [5] Stach, K., Sakamoto, S. Baba, K. Development of an ultrasound system for 3d-reconstruction of the foetus. volume 17, 1989.
- [6] M. Bajura, Henry Fuchs, and R. Ohbuchi. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Comput. Graphics*, 26(2):203–210, 1992.
- [7] W. Barrett and E. Bess. Interpolation by directed distance morphing. In R. A. Robb, editor, *Visualization in Biomedical Computing 1994, Proc. SPIE 2359*, pages 110–121, Rochester, MN, 1994.
- [8] J. Beier. *Automatische Quantifizierung von Koronarstenosen aus angiographischen Röntgenbildern, Fortschr.-Ber. VDI Reihe 17 Nr.95*. VDI-Verlag, Düsseldorf, 1993.
- [9] James F. Blinn. Light reflection functions for simulation of clouds and dusty surfaces. *Comput. Graphics*, 16(3):21–29, 1982.
- [10] J.D. Boissinat. Surface Reconstruction from Planar Cross-Sections. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 393 – 397, June 1985.
- [11] Michael Bomans. Segmentationsverfahren zur 3D-Visualisierung von Kernspintomogrammen des Kopfes: Evaluierung der Standardverfahren und Entwurf und Realisierung eines klinisch einsetzbaren Systems. Dissertation, Fachbereich Informatik, Universität Hamburg, 1994.
- [12] Michael Bomans, Karl Heinz Höhne, Ulf Tiede, and Martin Riemer. 3D-segmentation of MR-images of the head for 3D-display. *IEEE Trans. Med. Imaging*, MI-9(2):177–183, 1990.
- [13] M. E. Brummer, R. M. Mersereau, R. L. Eisner, and R. R. J. Lewine. Automatic detection of brain contours in MRI data sets. *IEEE Trans. Med. Imaging*, 12:153–166, 1993.
- [14] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intell.*, PAMI-8(6):679–698, 1985.
- [15] L. S. Chen, Gabor T. Herman, R. A. Reynolds, and Jayaram K. Udupa. Surface shading in the cuberille environment. *IEEE Comput. Graphics Appl.*, 5(12):33–43, 1985.
- [16] Harvey E. Cline, William E. Lorensen, Ron Kikinis, and Ferenc Jolesz. Three-dimensional segmentation of MR images of the head using probability and connectivity. *J. Comput. Assist. Tomogr.*, 14(6):1037–1045, 1990.
- [17] Harvey E. Cline, William E. Lorensen, S. Ludke, C. R. Crawford, and B. C. Teeter. Two algorithms for three-dimensional reconstruction of tomograms. *Med. Phys.*, 15(3):320–327, 1988.
- [18] Harvey E. Cline, William E. Lorensen, S. P. Souza, Ferenc A. Jolesz, Ron Kikinis, Guido Gerig, and T. E. Kennedy. 3D surface rendered MR images of the brain and its vasculature. *J. Comput. Assist. Tomogr.*, 15(2):344–351, 1991.
- [19] D. J. David, D. C. Hemmy, and R. D. Cooter. *Craniofacial Deformities: Atlas of Three-Dimensional Reconstruction from Computed Tomography*. Springer-Verlag, New York, 1990.
- [20] Shotton D.M. Confocal scanning optical microscopy and its applications for biological specimens. volume 94, pages 175–206, 1989.
- [21] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. *Comput. Graphics*, 22(4):65–74, 1988.

- [22] Robert A. Drebin, Donna Magid, D. D. Robertson, and Elliot K. Fishman. Fidelity of three-dimensional CT imaging for detecting fracture gaps. *J. Comput. Assist. Tomogr.*, 13(3):487–489, 1989.
- [23] Hans-Heino Ehrlicke and G. Laub. Combined 3D-display of cerebral vasculature and neuroanatomic structures in MRI. In Karl Heinz Höhne, Henry Fuchs, and Stephen M. Pizer, editors, *3D-Imaging in Medicine: Algorithms, Systems, Applications*, volume 60 of *NATO ASI Series F*, pages 229–239. Springer-Verlag, Berlin, 1990.
- [24] L.A. Feldkamp, L.C.Davis, J.W. Kress. Practical Cone-Beam Algorithm. *Journal of the Optical Society of Amerika A*, 1(6):612–619, 1989.
- [25] Elliot K. Fishman, Derek R. Ney, and Donna Magid. Three-dimensional imaging: Clinical applications in orthopedics. In Karl Heinz Höhne, Henry Fuchs, and Stephen M. Pizer, editors, *3D-Imaging in Medicine: Algorithms, Systems, Applications*, volume 60 of *NATO ASI Series F*, pages 425–440. Springer-Verlag, Berlin, 1990.
- [26] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley Publ. Comp., Reading, MA, 2. edition, 1990.
- [27] Bernhard Geiger and Ron Kikinis. Simulation of endoscopy. In Nicholas Ayache, editor, *Computer Vision, Virtual Reality and Robotics in Medicine, Proc. CVRMed '95*, volume 905 of *Lecture Notes in Computer Science*, pages 277–281. Springer-Verlag, Berlin, 1995.
- [28] Guido Gerig, O. Kübler, Ron Kikinis, and Ferenc A. Jolesz. Nonlinear anisotropic filtering of MRI data. *IEEE Trans. Med. Imaging*, 11(2):221–232, 1992.
- [29] Guido Gerig, J. Martin, Ron Kikinis, O. Kübler, M. Shenton, and Ferenc A. Jolesz. Automating segmentation of dual-echo MR head data. In Alan C. F. Colchester and D.J. Hawkes, editors, *Information Processing in Medical Imaging, Proc. IPMI '91*, volume 511 of *Lecture Notes in Computer Science*, pages 175–187. Springer-Verlag, Berlin, 1991.
- [30] S. Großkopf, A. Hildebrand. Three-Dimensional Reconstruction of Coronary Arteries from X-Ray Projections. In P. Lanzer, M. Lipton, editor, To be published: *Vascular Diagnostics: Principles and Technology*. Springer-Verlag, Heidelberg, 1996.
- [31] S. Haring, Max A. Viergever, and J. N. Kok. A multiscale approach to image segmentation using Kohonen networks. In H. H. Barrett and A. F. Gmitro, editors, *Information Processing in Medical Imaging, Proc. IPMI '93*, volume 687 of *Lecture Notes in Computer Science*, pages 212–224. Springer-Verlag, Berlin, 1993.
- [32] D. C. Hemmy and P. L. Tessier. CT of dry skulls with craniofacial deformities: Accuracy of three-dimensional reconstruction. *Radiology*, 157(1):113–116, 1985.
- [33] A. Hildebrand. Bestimmung computer-graphischer Beschreibungsattribute für reale 3D-Objekte mittels Analyse von 2D-Rasterbildern. PhD Thesis, TH Darmstadt, Darmstadt, 1996.
- [34] A. Hildebrand and S. Großkopf. 3D Reconstruction of Coronary Arteries from X-Ray Projections. In *Proceedings of the Computer Assisted Radiology CAR'95 Conference*, Berlin, 1995. Springer-Verlag.
- [35] W. Hiltman. Die 3d-strukturkonstruktion aus ultraschallbildern. October 1994.
- [36] Karl Heinz Höhne and Ralph Bernstein. Shading 3D-images from CT using gray level gradients. *IEEE Trans. Med. Imaging*, MI-5(1):45–47, 1986.
- [37] Karl Heinz Höhne, Michael Bomans, Bernhard Pflessner, Andreas Pommert, Martin Riemer, Thomas Schiemann, and Ulf Tiede. Anatomic realism comes to diagnostic imaging. *Diagn. Imaging*, (1):115–121, 1992.
- [38] Karl Heinz Höhne, Michael Bomans, Andreas Pommert, Martin Riemer, Carsten Schiers, Ulf Tiede, and Gunnar Wiebecke. 3D-visualization of tomographic volume data using the generalized voxel-model. *Visual Comput.*, 6(1):28–36, 1990.
- [39] Karl Heinz Höhne and W. A. Hanson. Interactive 3D-segmentation of MRI and CT volumes using morphological operations. *J. Comput. Assist. Tomogr.*, 16(2):285–294, 1992.
- [40] Karl Heinz Höhne, Bernhard Pflessner, Andreas Pommert, Martin Riemer, Thomas Schiemann, Rainer Schubert, and Ulf Tiede. A new representation of knowledge concerning human anatomy and function. *Nature Med.*, 1(6):506–511, 1995.
- [41] A. Höss, J. Debus, R. Bendl, R. Engenhardt-Cabillic, and Wolfgang Schlegel. Computerverfahren in der dreidimensionalen Strahlentherapieplanung. *Radiologe*, 35(9):583–586, 1995.
- [42] J. Jäger. Volumetric Reconstruction of Heart Ventricles from X-Ray Projections, 1996. diploma thesis, in german.

- [43] M. Kass, A. Witkin, D. Terzopoulos. Snakes: Active Contour Models. *IEEE First International Conference on Computer Vision*, pages 259–268, 1987.
- [44] Arie Kaufman, editor. *Volume Visualization*. IEEE Computer Society Press, Los Alamitos, CA, 1991.
- [45] Arie Kaufman, Karl Heinz Höhne, Wolfgang Krüger, Larry J. Rosenblum, and Peter Schröder. Research issues in volume visualization. *IEEE Comput. Graphics Appl.*, 14(2):63–67, 1994.
- [46] M. L. Kessler and D. L. McShan. An application for design and simulation of conformal radiation therapy. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1994, Proc. SPIE 2359*, pages 474–483, Rochester, MN, 1994.
- [47] Ron Kikinis, Martha E. Shenton, Dan V. Iosifescu, Robert W. McCarley, Pairash Saiviroonporn, Hiroto H. Hokama, Andre Robotino, David Metcalf, Cynthia G. Wible, Chiara M. Portas, Robert M. Donnino, and Ferenc A. Jolesz. A digital brain atlas for surgical planning, model driven segmentation, and teaching. *IEEE Trans. Visualization Comput. Graphics*, 2(3):232–241, 1996.
- [48] Teuvo Kohonen. *Self-Organisation and Associative Memory*. Springer-Verlag, Berlin, 2. edition, 1988.
- [49] Wolfgang Krueger and Bernd Froehlich. The responsive workbench. *IEEE Comput. Graphics Appl.*, 14(3):12–15, 1994.
- [50] D. Laur and Pat Hanrahan. Hierarchical splatting: A progressive refinement algorithm for volume rendering. *Comput. Graphics*, 25(4):285–288, 1991.
- [51] J. Leugyel, D.P. Greenberg, and Popp R. Time-Dependent Three-Dimensional Intravascular Ultrasound. In *Computer Graphics Proceedings*, pages 457–464, Los Angeles, 1995. SIGGRAPH.
- [52] M. Levoy. Display of Surface from Volume Data. *IEEE Computer Graphics and Applications*, 8(3):29–37, 1988.
- [53] Marc Levoy. Display of surfaces from volume data. *IEEE Comput. Graphics Appl.*, 8(3):29–37, 1988.
- [54] Marc Levoy. A hybrid ray tracer for rendering polygon and volume data. *IEEE Comput. Graphics Appl.*, 10(2):33–40, 1990.
- [55] Lun-Jou Lo, Jeffrey L. Marsh, Michael W. Vannier, and Vikas V. Patel. Craniofacial computer-assisted surgical planning and simulation. *Clin. Plast. Surg.*, 21(4):501–516, 1994.
- [56] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *Comput. Graphics*, 21(4):163–169, 1987.
- [57] Maria Magnusson, Reiner Lenz, and Per-Erik Danielsson. Evaluation of methods for shaded surface display of CT volumes. *Comput. Med. Imaging Graph.*, 15(4):247–256, 1991.
- [58] D. Marr and E. Hildreth. Theory of edge detection. *Proc. R. Soc. Lond.*, B 207:187–217, 1980.
- [59] Stephen R. Marschner and Richard J. Lobb. An evaluation of reconstruction filters for volume rendering. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proc. Visualization '94*, pages 100–107, Los Alamitos, CA, 1994. IEEE Computer Society Press.
- [60] Wido Menhardt. Iconic fuzzy sets for MR image segmentation. In A. E. Todd-Pokropek and Max A. Viergever, editors, *Medical Images: Formation, Handling and Evaluation*, volume 98 of *NATO ASI Series F*, pages 579–591. Springer-Verlag, Berlin, 1992.
- [61] R. Millner. *Ultraschalltechnik, Grundlagen und Anwendungen*.
- [62] B. S. Morse, Stephen M. Pizer, and A. Liu. Multiscale medial analysis of medical images. *Image Vision Comput.*, 12(6):327–338, 1994.
- [63] Shigeru Muraki. Volume data and wavelet transforms. *IEEE Comput. Graphics Appl.*, 13(4):50–56, 1993.
- [64] Shigeru Muraki. Multiscale volume representation by a DOG wavelet. *IEEE Trans. Visualization Comput. Graphics*, 1(2):109–116, 1995.
- [65] P.J. Neugebauer. Interactive Segmentation of Dentistry Range Images in CIM Systems for the Construction of Ceramic Inlays using Edge Tracing. In *Proceedings of the Computer Assisted Radiology CAR'95 Conference*, Berlin, 1995. Springer-Verlag.
- [66] Derek Ney and Elliot K. Fishman. Editing tools for 3D medical imaging. *IEEE Comput. Graphics Appl.*, 11(6):63–70, 1991.

- [67] Derek Ney, Elliot K. Fishman, Donna Magid, D. D. Robinson, and A. Kawashima. Three-dimensional volumetric display of CT data: Effect of scan parameters upon image quality. *J. Comput. Assist. Tomogr.*, 15(5):875–885, 1991.
- [68] P. Ning and J. Bloomenthal. An evaluation of implicit surface tilers. *IEEE Comput. Graphics Appl.*, 13(6):33–41, 1993.
- [69] Bernhard Pflessner, Ulf Tiede, and Karl Heinz Höhne. Towards realistic visualization for surgery rehearsal. In N. Ayache, editor, *Computer Vision, Virtual Reality and Robotics in Medicine, Proc. CVRMed '95*, volume 905 of *Lecture Notes in Computer Science*, pages 487–491. Springer-Verlag, Berlin, 1995.
- [70] Andreas Pommert, Michael Bomans, and Karl Heinz Höhne. Volume visualization in magnetic resonance angiography. *IEEE Comput. Graphics Appl.*, 12(5):12–13, 1992.
- [71] Andreas Pommert, Wolf-Joachim Hölzje, Nikolaus Holzknicht, Ulf Tiede, and Karl Heinz Höhne. Accuracy of images and measurements in 3D bone imaging. In Heinz U. Lemke, Michael L. Rhodes, C. Carl Jaffe, and Roland Felix, editors, *Computer Assisted Radiology, Proc. CAR '91*, pages 209–215. Springer-Verlag, Berlin, 1991.
- [72] Andreas Pommert, Rainer Schubert, Martin Riemer, Thomas Schiemann, Ulf Tiede, and Karl Heinz Höhne. Symbolic modeling of human anatomy for visualization and simulation. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1994, Proc. SPIE 2359*, pages 412–423, Rochester, MN, 1994.
- [73] Andreas Pommert, Ulf Tiede, Gunnar Wiebecke, and Karl Heinz Höhne. Surface shading in tomographic volume visualization: A comparative study. In *First Conference on Visualization in Biomedical Computing, Proc. VBC '90*, pages 19–26. IEEE Computer Society Press, Los Alamitos, CA, 1990.
- [74] S. P. Raya and Jayaram K. Udupa. Low-level segmentation of 3-D magnetic resonance brain images: A rule-based system. *IEEE Trans. Med. Imaging*, MI-9(3):327–337, 1990.
- [75] A. Rosenfeld and A. C. Kak. *Digital Picture Processing*. Academic Press, New York, 1982.
- [76] H. Rusinek, M. E. Noz, G. Q. Maguire, A. Kalvin, B. Haddad, D. Dean, and C. Cutting. Quantitative and qualitative comparison of volumetric and surface rendering techniques. *IEEE Trans. Nucl. Sci.*, 38(2):659–662, 1991.
- [77] A. Hildebrand, R. Malkewitz, W. Mueller, R. Ziegler, G. Graschew S. Grosskopf. Computer aided surgery – vision and feasibility of an advanced operation theatre. volume 20, pages 825–835. Pergamon, November 1996.
- [78] Bono, P. Sakas, G. Special issue on medical visualization. volume 20, pages 759–838. Pergamon, November 1996.
- [79] Schreyer, L., Grimm, M. Sakas, G. Case study: Visualization of 3d-ultrasonic data. pages 369–373. IEEE, October 1994.
- [80] Vicker, M., Plath., P. Sakas, G. Case study: Visualization of laser confocal microscopy data. pages 375–380. IEEE, October 1996.
- [81] Walter, S. Sakas, G. Extracting surfaces from fuzzy 3d-ultrasonic data. pages 465–474. Addison Wesley, August 1995.
- [82] Thomas Schiemann, Michael Bomans, Ulf Tiede, and Karl Heinz Höhne. Interactive 3D-segmentation. In Richard A. Robb, editor, *Visualization in Biomedical Computing II, Proc. SPIE 1808*, pages 376–383, Chapel Hill, NC, 1992.
- [83] Thomas Schiemann, Bernhard Dippold, Rainer Schmidt, Andreas Pommert, Martin Riemer, Rainer Schubert, Ulf Tiede, and Karl Heinz Höhne. 3D visualization for radiotherapy treatment planning. In Heinz U. Lemke, Michael L. Rhodes, C. Carl Jaffe, and Roland Felix, editors, *Computer Assisted Radiology, Proc. CAR '93*, pages 669–675. Springer-Verlag, Berlin, 1993.
- [84] Thomas Schiemann, Karl Heinz Höhne, Christoph Koch, Andreas Pommert, Martin Riemer, Rainer Schubert, and Ulf Tiede. Interpretation of tomographic images using automatic atlas lookup. In Richard A. Robb, editor, *Visualization in Biomedical Computing 1994, Proc. SPIE 2359*, pages 457–465, Rochester, MN, 1994.
- [85] Thomas Schiemann, Ulf Tiede, and Karl Heinz Höhne. Segmentation of the Visible Human for high quality volume based visualization. *Med. Image Anal.*, 1(4), 1997. (in press).
- [86] Rainer Schmidt, Thomas Schiemann, Wolfgang Schlegel, Karl Heinz Höhne, and K.-H. Hübener. Consideration of time-dose-patterns in 3D treatment planning: An approach towards 4D treatment planning. *Strahlenther. Onkol.*, 170(5):292–301, 1994.

- [87] W. J. Schroeder, J. A. Zarge, and William E. Lorensen. Decimation of triangle meshes. *Comput. Graphics*, 26(2):65–70, 1992.
- [88] Rainer Schubert, Karl Heinz Höhne, Andreas Pommert, Martin Riemer, Thomas Schiemann, and Ulf Tiede. Spatial knowledge representation for visualization of human anatomy and function. In Harrison H. Barrett and A. F. Gmitro, editors, *Information Processing in Medical Imaging, Proc. IPMI '93*, volume 687 of *Lecture Notes in Computer Science*, pages 168–181. Springer-Verlag, Berlin, 1993.
- [89] Rainer Schubert, Wolf-Joachim Höltje, Ulf Tiede, and Karl Heinz Höhne. 3D-Darstellungen für die Kiefer- und Gesichtschirurgie. *Radiologe*, 31:467–473, 1991.
- [90] Victor Spitzer, Michael J. Ackerman, Ann L. Scherzinger, and David Whitlock. The Visible Human male: A technical report. *J. Am. Med. Inf. Ass.*, 3(2):118–130, 1996.
- [91] Colin Studholme, Derek L. G. Hill, and David J. Hawkes. Automated 3-D registration of MR and CT images of the head. *Med. Image Anal.*, 1(2):163–175, 1996.
- [92] Russell H. Taylor, Stephane Lavallée, Grigore C. Burdea, and Ralph Mösges. *Computer Integrated Surgery: Technology and Clinical Applications*. MIT Press, Cambridge, MA, 1995.
- [93] Ulf Tiede, Karl Heinz Höhne, Michael Bomans, Andreas Pommert, Martin Riemer, and Gunnar Wiebecke. Investigation of medical 3D-rendering algorithms. *IEEE Comput. Graphics Appl.*, 10(2):41–53, 1990.
- [94] Ulf Tiede, Thomas Schiemann, and Karl Heinz Höhne. Visualizing the Visible Human. *IEEE Comput. Graphics Appl.*, 16(1):7–9, 1996.
- [95] Arthur W. Toga and John C. Mazziotta. *Brain Mapping*. Academic Press, San Diego, CA, 1996.
- [96] T. Totsuka and Marc Levoy. Frequency domain volume rendering. *Comput. Graphics*, pages 271–278, 1993.
- [97] Petra A. van den Elsen, E.-J. D. Pol, and Max A. Viergever. Medical image matching: A review with classification. *IEEE Engng. Med. Biol. Magazine*, 12(1):26–39, 1993.
- [98] Michael W. Vannier, C. F. Hildebolt, J. L. Marsh, T. K. Pilgram, W. H. McAlister, G. D. Shackelford, C. J. Offutt, and R. H. Knapp. Craniosynostosis: Diagnostic value of three-dimensional CT reconstruction. *Radiology*, 173:669–673, 1989.
- [99] Ellen Vaske. Segmentation von Kernspintomogrammen mit der topologischen Karte zur 3D-Visualisierung. IMDM Institutsbericht 91/1, Institut für Mathematik und Datenverarbeitung in der Medizin, Universität Hamburg, 1991.
- [100] G. Vosselman. *Relational Matching*. Springer-Verlag, 1992.
- [101] A. Wahle, E. Wellnhofer, I. Mugaragu, A. Trebeljahr, H. Oswald, E. Fleck. Application of Accurate 3D Reconstruction from Biplane Angiograms in Morphometric Analyses and in Assessment of Diffuse Coronary Artery Disease. In *CAR'95: Computer Assisted Radiology*, Berlin, Germany, June 21 - 24 1995. Springer Verlag.
- [102] Å. Wallin. Constructing isosurfaces from CT data. *IEEE Comput. Graphics Appl.*, 11(6):28–33, 1991.
- [103] Suzanne J. Weghorst, Hans B. Sieburg, and Karen S. Morgan, editors. *Medicine Meets Virtual Reality: Health Care in the Information Age, Proc. MMVR '96*, volume 29 of *Studies in Health Technology and Informatics*. IOS Press, Amsterdam, 1996.
- [104] William M. Wells III, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis. Multi-modal volume registration by maximization of mutual information. *Med. Image Anal.*, 1(1):35–51, 1996.
- [105] Lee Westover. Footprint evaluation for volume rendering. *Comput. Graphics*, 24(4):367–376, 1990.
- [106] J. Wilhelms and A. van Gelder. Topological considerations in isosurface generation. *Comput. Graphics*, 24(5):79–86, 1990.
- [107] Frank Wilmer, Ulf Tiede, and Karl Heinz Höhne. Reduktion der Oberflächenbeschreibung triangulierter Oberflächen durch Anpassung an die Objektform. In S. Fuchs and R. Hoffmann, editors, *Musternererkennung 1992, Proc. 14. DAGM-Symposium*, pages 430–436. Springer-Verlag, Berlin, 1992.
- [108] P. H. Winston. *Artificial Intelligence*. Addison-Wesley, Reading, MA, 3. edition, 1992.
- [109] Roni Yagel, D. Cohen, and Arie Kaufman. Discrete ray tracing. *IEEE Comput. Graphics Appl.*, 12(5):19–28, 1992.

- [110] Takami Yasuda, Y. Hashimoto, Shigeki Yokoi, and Jun-Ichiro Toriwaki. Computer system for craniofacial surgical planning based on CT images. *IEEE Trans. Med. Imaging*, MI-9(3):270–280, 1990.
- [111] Frans W. Zonneveld and Keizo Fukuta. A decade of clinical three-dimensional imaging: a review. part 2: Clinical applications. *Invest. Radiol.*, 29:574–589, 1994.
- [112] Frans W. Zonneveld, S. Lobregt, J. C. H. van der Meulen, and J. M. Vaandrager. Three-dimensional imaging in craniofacial surgery. *World J. Surg.*, 13:328–342, 1989.