VOLUMETRIC RENDERING OF MEDICAL DATA:
Applications to stereotactic neurosurgery planning

by

Donald Louis Collins

Department of Electrical Engineering
McGill University
Montréal, Québec

A Thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Master of Engineering.

Abstract

This thesis addresses the problem of registration, merging and viewing medical images from multiple modalities for stereotactic surgery planning. A technique is presented, using *volumetric rendering* of tomographic data, to create anatomical perspective projections that can be easily merged with vascular projections.

The geometry of the digital subtraction angiogram (DSA) projection is reproduced by the rendering process to enable registration between the DSA radiograph and the volumetric projection. The viewing parameters are calculated from the location of fiducial markers in the image. Each angiogram is overlayed onto a translucent volumetric projection of computed tomography (CT) or magnetric resonance (MR) data, rendered to give a matched view.

Until recently, rendered medical volumes have been used qualitatively in the diagnostic and surgical planning process. Since the volumetric projections are matched to the DSA images, the complete set of stereotactic surgery planning tools can be used to identify points and measure distances in the rendered images.

A point spread function of the rendering process is derived to establish a theoretical limit on the accuracy of the technique and is verified by experimentation.

Résumé

Cette thèse étudie le problème de la corrélation spatiale, de la surimposition et de la visualisation d'images médicales de modalitées multiples pour la planification de chirugie stéréotaxique. Le but de cette thèse est de développer une technique utilisant la projection volumétrique d'images tomographiques, afin de créér des projections d'anatomies en perspective sur lesquelles des angiogrammes numérisés soustraits (ANS) peuvent être surimposés.

Afin que la surimposition soit alignée, la géométrie de la projection ANS est reproduite par le processus de projection volumétrique. Les paramètres de visualization sont calculés à partir de la position de pointes fiduciaires dans l'image. Chaque angiogramme est surimposé sur une projection volumétrique translucide de données tomographiques ou d'imagerie par résonance magnétique, créée pour avoir la même vue.

Jusqu'à maintenant, les images médicales volumétriques sont utilisées de façon qualitative seulement pour le diagnostique et la planification de la chirugie. Puisque les projections volumétriques sont en corrélation avec les angiogrammes, un ensemble d'outils informatisés pour faire la planification de la chirugie stéréotaxique peuvent être utilisés pour indentifier des points et mesurer des distances dans les images volumétriques.

La function de difraction ('point-spread') de la projection volumétrique est dérivée mathématiquement et verifiée expérimentalement pour établir une limite théorique sur la précision de la technique.

Acknowledgements

I wish to express my sincere thanks to Dr. T.M. Peters, my advisor, for direction, encouragement and making sure that I get it write. As co-founder of the NeuroImaging Lab, he has made possible the environment necessary for this research.

My gratitude is extended to all my associates at the Montreal Neurological Institute who have helped at one time or another. My thanks goes out to Bruce Pike and Dr. Alan Evans for enlightening discussions, to Chris Henri for his indispensible collaboration to develop the merging algorithm and finally, to Sean Marrett for unlimited assistance and patient explanation of things technical, for friendship and moral support.

I would like to take this opportunity to thank the 'Fonds pour la Formation de Chercheurs et l'Aide à la Recherche' (FCAR) for the financial support made available by a graduate level bursary. I am also grateful for hardware support provided by Philips Medical Systems.

to my parents

Contents

List of Figures

List of Symbols

symbol	**meaning**
$\mathcal{O}$	real physical object
(x, y, z)	point in world (frame) coordinate system
(X, Y, Z)	point volume coordinate system
(u, v, w)	point screen (image plane) coordinate system
$\mathcal{O}_f$	object after lowpass filtering
$\rho()$	low-pass object density function
$\rho_s()$	sampled low-pass object density function
$\mathcal{V}$	volumetric data set (volume)
$h_v()$	volumetric reconstruction kernel
V	voxel
$\mathcal{P}()$	ideal projection operator
$P()$	volume rendered projection operator
x_s, y_s, z_s	volume sampling intervals in the $x-$, $y-$ and $z-$ directions, respectively
$comb()$	infinite 3D array of unit response functions
$sinc()$	$sinc(x) = sin(\pi x)/\pi x$
f_x, f_y, f_z	object spatial frequency in the $x-$, $y-$ and $z-$ directions, respectively

symbol	meaning
B_x, B_y, B_z	spatial frequency limits in the $x-$, $y-$ and $z-$ directions, respectively
$\mathbf{i}$	a sample point (i_x, i_y, i_z) falling within the extent of the reconstruction kernel $h_v()$ centered on (x, y, z)
$I(u, v)$	intensity of pixel (u,v) in image plane
T_{wv}	world to voxel coordinate system transformation
T_{vs}	voxel to screen coordinate system transformation
O_c	object contrast
I_c	image contrast
S_{max}	maximum object intensity
S_{min}	minimum object intensity
I_{max}	maximum image intensity
I_{min}	minimum image intensity
PSF	point spread function
MTF	modulation transfer function
LSF	line spread function
$\mathbf{S}$	scaling matrix
$\mathbf{P}$	perspective scaling matrix
$\mathbf{C}$	cropping matrix

symbol	meaning
$\mathcal{P}$	piercing point
$\mathcal{L}$	center of projection
θ	rotation about $x-$axis
ϕ	rotation about $y-$axis
ψ	rotation about $z-$axis
$\mathcal{R}$	viewing ray
t_x, t_y, t_z	translations in the $x-$, $y-$ and $z-$ directions, respectively
s_x, s_y, s_z	scaling in the $x-$, $y-$ and $z-$ directions, respectively
f	focal distance
k_u, k_z	imaging scaling in the $u-$ and $v-$directions
$\tau_{\mathbf{i}}$	contribution of sample **i** to the volumetric reconstruction

Chapter 1

Introduction

Stereotactic neurosurgery is a technique used to safely approach lesions deep within the brain otherwise impossible to reach using traditional methods [49,6,58]. Two constraints must be met when planning the path of a probe or needle to a cerebral target. The path must not pass through important cortical structures unless absolutely necessary, nor must it injure a blood vessel. In this thesis, a technique is presented to merge the digital subtraction angiogram (DSA) projection with the tomographic images in stereo. Viewing the three-dimensional (3D) cerebral vasculature in conjunction with the surrounding anatomy when planning the path of a probe will lead to safer stereotactic neurosurgery.

Invasive stereotactic procedures are carried out through a small burr hole in the skull and include biopsy, aspiration of lesions, implantation of electrodes for the recording of deep EEG signals [75] and deposition of radioactive seeds in tumours. Non-invasive radio-surgical techniques use an external beam of high energy ionizing radiation from a linear accelerator to treat an affected area [78,77]. For all of these procedures, precise localization of the target point with respect to some fixed external reference frame is essential.

Stereotaxis literally means *spatial organization*. Fixation of a rigid reference frame to the patient's skull during imaging and surgery establishes an accurate 3D coordinate system for the cerebral volume (see Fig. 1.1). During scanning, fiducial marker plates

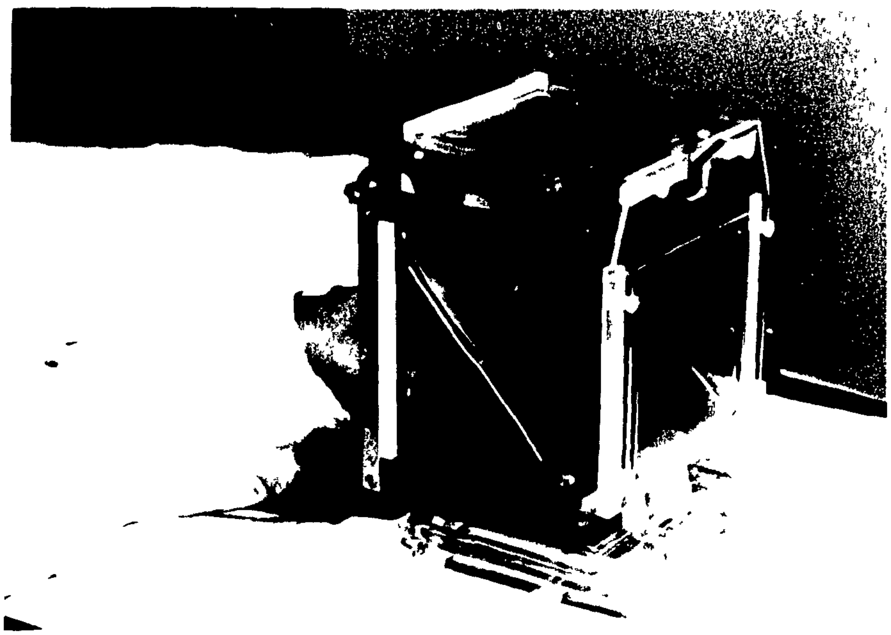

Figure 1.1: Stereotactic frame fixed to patient before imaging.

attached to the frame appear in each image, providing identifiable landmarks that define the imaging geometry. After scanning, the frame serves as a base for specialized instruments during surgery. Thus the frame enables the location of any point appearing in an image to be accurately determined and used as a target in surgery [49,67].

Medical imaging tools present a wealth of information to be used analytically and diagnostically. Vascular structure is visualized in digital subtraction angiograms (DSA) [74], most of which are performed in stereo at the Montreal Neurological Institute (MNI) [102]. The perception of depth due to the stereo cues is found to aid qualitatively in the diagnostic and surgical planning process [65]; however, the relationship between the vasculature and the surrounding tissue is difficult to establish due to the lack of anatomical information resulting from the subtraction. New tomographic techniques provide anatomical cross-sectional images in which organs are seen clearly, unobscured by others, permitting a radiologist or surgeon to locate points of interest deep within an organ. X-ray computed tomography (CT) provides detail of the structural tissues while magnetic resonance (MR) images yield high-contrast in soft tissues.

Examination and diagnosis is facilitated if the complementary information from each modality is overlayed and compared [20]. While this is a fairly trivial task for CT and MR (logistical constraints aside), merging a DSA perspective projection image with tomographic MR and CT images is more difficult since their native geometric formats differ greatly. This thesis describes a technique, using *volumetric rendering* of tomographic data, to create anatomical perspective projections that may easily be merged with vascular projections. (Note on terminology: *Volume rendering* is used to describe display methods that use overpainting of voxels to form a projection image in the field of image processing. *Volume compositing* is the term used in computer graphics to describe techniques that are based on the blending of semi-transparent voxels. Although both terms are used interchangably, the work in this thesis follows the latter definition [52].)

Most software systems used for medical image analysis and surgical planning are based on a radiographic view-box model and as such are limited to the use of single two-dimensional (2D) tomographic and projection image. It is a formidable task to appreciate the intricate spatial relationships between different anatomical structures surrounding a target since the planar format limits 3D spatial comprehension. Mentally reconstructing the unknown shape of a tumour or lesion is difficult. Reorganization of the data by reslicing or reprojection is useful, but the observer must attempt to visualize the 3D form.

The goal of 3D reconstruction by computer rendering is to accurately and non-invasively display an organ (or organs) as seen inside the body. The additional information, both qualitative and quantitative, and the supplementary insight gleaned from the 3D model without the associated risks of explorative surgery makes this approach desirable.

The appearance of dedicated turnkey systems for 3D volumetric rendering of CT and MR data from a number of the major vendors shows a definite interest in the field[1]. Neurosurgeons and radiologists are demanding this type of technology[2], since visualization of anatomical structures in three dimensions is much more natural than examining cross sectional data.

Until the mid-seventies 3D reconstruction was difficult and time-consuming. The contour of the object of interest was outlined by hand on the tomographic slices and reconstructions took the form of 3D models made of wood, plexiglass or cardboard [61]. They were costly, construction time was lengthly and it was difficult to obtain any quantitative measurements such as shape or surface area. New models were required if sections of the object of interest were to be examined.

Computer graphic systems have solved many of these difficulties. The contours are created automatically by computer using an edge tracking scheme on a threshold image and then displayed from any angle on a vector graphics terminal using simple 3D transformations [61]. A surface corresponding to these planar outlines are reconstructed using simple triangulation to define a series of ribbon strips between each two consecutive contours [44,24]. This 3D wireframe mesh is viewed using known geometric algorithms to depict perspective with or without hidden-line removal. The triangulated models are displayed on raster graphics terminals with more complex algorithms to shade the polygonal surfaces using Phong or Gouraud shading [76,32], greatly enhancing the apparent quality of the image. In addition to the visualization techniques described above, algorithms were also developed to estimate surface area and volume [15].

The cuberille models developed by Herman and others [38,31,84,40,13,97] visualize the volumetric data by projecting the sides of each voxel[3] lying on the surface of interest directly on to the imaging plane. These techniques have been used clinically in a number of institutions for craniofacial surgery [100], reconstructive surgery planning [43,91], the

[1]ISG Technology's CAMRA system, Siemen's 3D Display; Philip's MSP Package; General Electric's 3D Package; Hipgraphic's Orthotool.

[2]Personal communication with Drs. A. Olivier and R. Ethier.

[3]Voxels are the 3D equivalent of image pixels.

design of customized artificial joints and prosthesis [87,96] and the visualization of bone fractures [86,22,59,9,60].

The implicit assumption with traditional computer-graphics based methods is that the 3D volume may be accurately represented by a model of thin surfaces suspended in a transparent environment such as air [90]. For 'hard-edged' objects, these techniques suffice, but there are many types of data where this scheme is not applicable. Where the data is intrinsically ill-defined, such as in positron emission tomography (PET), the boundary between regions is not a 3D surface but a 3D transition zone. New techniques are needed for the visualization of these 'soft-edged' objects.

3D reconstruction by direct volume compositing uses computer graphics techniques to project all of the pixels of the tomographic slices (all of the voxels from the volume) onto a 2D plane without surface or edge approximations [54,53,25,16]. The final colour of an image pixel is dependent on the characteristics of all the voxels lying in front of it. Thus, 'soft' surfaces contribute to the final colour and are not lost by thresholding. The rendered image may be an orthographic or perspective projection, opaque or transparent, colour or grey-scale, shaded or surface enhanced, depending on the graphics routines used.

In this thesis, the anatomical tomographic data sets are rendered using these direct volume reconstruction techniques and they may be merged with the DSA projection image if both have the same projection characteristics. As discussed previously, exact registration of the radiographic and the volume rendered projections is dependent upon the use of the stereotactic frame. The geometry of the DSA projection is reproduced by the rendering process and is defined by viewing parameters determined by the location of the fiducial markers in the image [92]. The tomographic volume is rendered in perspective as if seen from the same viewing point as the angiogram [14,73].

To form a merged stereo pair, each angiogram is overlayed onto a translucent volumetric projection of CT or MR data that has been rendered to give a matched view. Using a computer equipped with a liquid-crystal polarizing screen, which displays left and right eye views of the data sequentially at 120 Hz, the operator has the ability to directly view the stereoscopic angiograms using circularly polarized glasses. The merged images may be viewed together as part of a single stereoscopic 3D image with the same system [36,35].

The following questions are addressed in this thesis:

- Are direct volume-rendered anatomical images useful for the planning of neurosurgery?
- Can these images be used quantitatively?
- Can the tomographic data sets be integrated with each other and with the radiographic projection images of digital subtraction angiography?

The goal of this work is not to develop a new rendering technique, but to revise and to extend the existing schemes. The original work in this thesis consists of the design of a method to integrate volumetric and projection data using existing volume rendering methods, the analysis of the quantitative accuracy of the rendering and merging procedures and the modification of an existing rendering algorithm to display medical data. Specifically, this is limited to the merging of medical tomographic data, such as MR and CT, with projection digital subtraction angiograms for use in neurosurgical planning. The rendering method is extended to correctly display the anatomical data present in the volume and create a precisely matched perspective projection image.

The procedure is based on the following assumptions, namely that an external fiducial structure is used to provide a common frame of reference in the images between the modalities [67], and the projection images are artifact-free. This is the case at the MNI, although this may not be true in general [74]. It is believed however that these limitations may be overcome by addressing each in turn in a preprocessing operation.

The thesis begins with a description of the imaging methods used to collect medical data and the techniques available to establish a coordinate reference frame for measurement and analysis of structures within the images.

Chapter three commences with a definition of a volumetric data set and follows with a descriptive, critical overview of the the existing methods of 3D visualization.

Chapter four presents an existing rendering algorithm that overcomes many of the problems associated with the earlier schemes. Work performed to "tune" the application of this algorithm to the rendering of MR data in perspective is also described. Different techniques for 2D and 3D image registration are summarized; in particular, merging stereoscopic DSA with tomographic data is explained. This last section was performed in collaboration with Chris Henri, a fellow graduate student.

Chapter five presents the derivation of a point spread function for arbitrary perspective views of volume data sets. Since the tomographic volume is manipulated to yield a perspective rendering matched with the perspective projection of the angiogram, distortions of the image may arise when the tomographic data is scaled, translated and rotated. Experiments that yield estimates of the accuracy of the rendered images for potential 3D geometric measurement are described.

Chapter six summarizes the technique, discusses the applicability of using the 3D perspective rendered MR or CT volumes in conjunction with digital subtraction angiograms on a stereoscopic workstation for neurosurgical planning and provides direction for future work.

Chapter 2

Materials and methods

Stereotactic techniques used in neurosurgery and neurosurgical research are characterized by precise localization of cerebral targets in order to direct the tip of a needle, electrode or probe in 3D space to reach a specific point in the brain. The stereotactic patients at the MNI are imaged by a number of modalities routinely. DSA images are used in conjunction with X-ray CT and MRI to identify 3D targets within the cranial vault for analysis and planning of neurosurgery. This chapter briefly describes each of these imaging modalities, the stereotactic environment and finally the stereoscopic workstation used to analyse and plan the surgical techniques.

2.1 Imaging

2.1.1 Stereoscopic DSA

Digital subtraction angiography (DSA) is an invasive radiographic technique used for vasculature imaging. A catheter is inserted into an artery (or vein) and a radiopaque contrast material is injected into the blood flow. The technique involves the acquisition of X-ray images both before and during the flow of contrast media following the injection. A number of images are taken as the contrast bolus moves along with the blood through the main arteries, into the smaller branching arterioles and then into the capillaries

Figure 2.1: Stereoscopic digital subtraction angiogram. This stereo pair can be viewed using the crossed-eye technique so that the right eye fixes the image on the left and the left eye fixes the image on the right. When the two images are fused, the 3D structure is easily visualized.

followed by the small veins and finally into the main draining veins. Image processing of the angiograms is limited to subtraction of selected images as the bolus moves from the arterial phase to the venous phase. The subtraction of any two images eliminates the intervening tissues, leaving only the vessels where the contrast level differs between images [46]. A composite image is formed by subtracting the images that best show the arterial or venous phase of the vasculature from each other.

At the MNI, radiologists and surgeons view angiograms stereoscopically by using a stereoscope [45] or with a crossed eye technique (see Fig. 2.1). The use of stereo images provides depth information vital to assessing the spatial relationships between different vessels.

One of two methods is used to acquire stereoscopic angiograms on a Siemens Angiotron – Angioscop[1]. Either a dual-focus X-ray tube ([95]) is used or a C-arm gantry rotation is performed between the two required views. The 25mm focal spot separation of the stereoscopic tube gives a stereo disparity of 3 degrees at the image plane. Greater stereo shifts are acquired by angulating the gantry by approximately seven degrees between consecutive runs. In both cases, an image intensifier with variable field size is used, and 512 x 512 x 10 bit images are obtained from either a 19cm or 27cm field of view. Two mutually orthogonal image sequences are obtained; usually lateral and posterior-anterior (PA) projections.

Since the seven degree gantry angulation method provides a larger stereo disparity in the images and hence a greater perceivable depth, it is preferred over the stereoscopic X-ray tube. A new stereoscopic X-ray system will soon be installed at the MNI with a focal spot separation of 6.5 cm. This will give the benefits of both the stereoscopic tube and the wide disparity.

2.1.2 Computed tomography

X-ray computed tomography (CT) provides images that are accurate 2D maps of attenuation coefficients [37]. Contiguous axial slices are acquired with a GE 9800 scanner[2] with the stereotactic frame fixed to the scanner's couch, preventing patient movement. The CT images are reconstructed in a 512 x 512 matrix, 12 bits deep, with a field of view of 345 mm. If the gantry is angulated, oblique slices may be acquired. Look-up tables (LUTs) are applied to the CT data to enhance the visualization of soft-tissue contrast in the image.

[1]Siemens, Erlangen Germany.

[2]General Electric, Milwaukee.

2.1.3 Magnetic resonance

The MR image is based on the resonance behavior of hydrogen nuclei situated in a magnetic field, when stimulated by electro-magnetic energy at a radio-frequency. The contrast in the image is due to nuclear density as well as other tissue-related parameters that influence the MR signal, such as the relaxation properties (T1,T2) of the nuclei and the variation in resonant frequency of a particular nucleus with its molecular site (chemical shift) [82,47,70,71].

The MR image data is collected using a Philips 1.5T Gyroscan magnetic resonance imager using the body coil[3]. A 2D multi-slice spin-echo technique (TR 500ms, TE 30ms, 2 measurements, 3mm slice thickness) is used to gather contiguous axial or sagittal slices. Each slice is reconstructed in a 256 x 256 matrix, 12 bits deep, with a field of view of 325 mm. A 3D acquisition method is employed when numerous slices less than 2 mm thick are required, with slice thicknesses less than 3 mm giving the best results during the 3D rendering operation. The 3D spin echo has a better signal to noise ratio than the multislice technique. The 3D imaging parameters (TR 400ms, TE 30ms, 256x256x64 matrix, 1 measurement, 2mm slice thickness) are selected to keep imaging time at a minimum while obtaining a data volume large enough to cover the full frame and slices thin enough to achieve reasonable renderings.

2.2 Imaging computer

For the work described here, all volumetric image processing is performed on a PIXAR[4], a powerful imaging computer. Originally designed to composite computer generated special effects for science fiction films, its high quality imaging capabilities makes it ideal for the manipulation of medical data. The PIXAR consists of a large frame buffer, some fast local memory and a channel array processor (CHAP)[1]. Although limited by

[3]The stereotactic frame does not fit in the head coil.

[4]PIXAR, San Raphael California.

the lack of a floating point processor, software programs have been written using the ChapVolumes volumetric processing software library [2] to manipulate and render the tomographic data.

The author, in collaboration with Sean Marrett, has developed a software tool to register volumes of data [19]. The large frame buffer memory of the PIXAR allows two volumes to be stored in RAM (random access memory) simultaneously. The tool is used to translate, scale and rotate the volumes with respect to each other, and also to prepare a volume for a given projection.

Images are created by projecting each data point in the volume onto the image plane. High contrast is achieved by using the 36 bits per pixel available to define a large range of colours and shades without creating false contours.

2.3 Stereotactic environment

Stereotactic surgery is a technique whereby mechanical apparatus is used to direct and guide the approach of surgical instruments through a small burr hole in the skull towards small or deep-seated intra-cerebral lesions [6,58]. Stereotactic radiosurgery uses a similar apparatus to direct an external beam of high energy ionizing radiation to the target [78,49]. Surgical tasks include the section of deep nerve fiber tracts for the treatment of parkinsonism and intractable pain, implantation of radioactive seeds in tumors, evacuation of cerebral hemorrhage using stereotactically guided needles, electrolytic lesion, aspiration of abscesses and cysts, and biopsy of deep mass lesions. Such blind techniques differ from traditional surgery performed under direct visual control. Stereotactic procedures are the result of careful planning and reliance on specialized instruments. Before any of these procedures are attempted, the precise location of the "target" points (with respect to some fixed reference) within the brain must be accurately determined. Since 1985, an integrated environment has been created at the Montreal Neurological Institute and Hospital to meet this requirement [74].

Precise localization of points of interest is achieved using a rigid frame which is fixed to the patient's head with carbon-fiber or aluminum pins. Sets of fiducial markers are attached to the frame during imaging [67]. (An overview of stereotaxic technique is given in a report by Olivier in [66].) This section describes the frame and the markers used with the different modalities.

2.3.1 The frame

The frame used at the Montreal Neurological Institute is the "OBT" frame (Olivier, Bertrand and Tipal)[5]. Its design is based on, and it is physically similar to, the Leksell frame [50] although it embodies characteristics common to a number of other stereotactic frames as well [8,72,57]. Once attached to the patient's skull using carbon-fiber or aluminum pins, it stays fixed during imaging and surgery. Throughout the imaging session, the different modality-dependent fiducial marker plates are physically attached to the frame, while during surgery, it serves as a base for the stereotactic surgical instruments and maintains the coordinate reference system established during imaging.

The cube-shaped frame is designed to provide maximum strength, while being light enough for the patient to wear comfortably. Its base is constructed from an aluminum alloy on three sides and plastic (carbon fibre) on the fourth in order to avoid any closed loops which could produce spurious magnetic fields during MR imaging. The machined vertical posts accommodate a calibrated, horizontal, sliding bar which supports the stereotactic instrument carrier during surgery. The posts are engraved with millimeter scales so that the bar (which can be fixed to the vertical posts with thumb-screws) can be accurately positioned vertically. The carrier is equipped with a collar chuck and is mounted on the bar so that it can move horizontally.

When using the rectilinear system for electrode implantation or for biospy, the surgical tools are inserted directly into the carrier. A stereometer (also known as a phantom carrier) is used to determine the distance from the carrier stopper to any point within

[5]Fabricated by Tipal Instruments Inc, Montreal

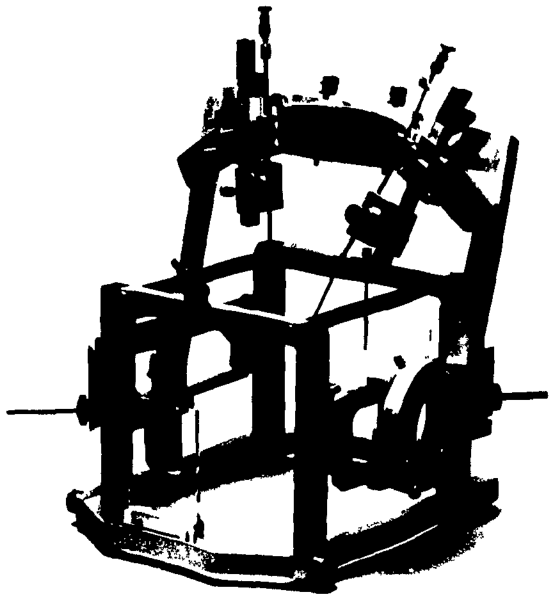

Figure 2.2: Stereotactic frame with movable semi-circular arc and instrument carrier.

the frame to gauge the insertion depth of the tool. Targets may be approached via a polar coordinate system using a semicircular arc mounted on the frame with a movable carrier (see Fig. 2.2). It is is fixed to carriers on either side of the frame so that its center coincides with the center of the target. Thus, any surgical tool placed in the carrier and inserted to a depth equal to the arc's radius will reach the target in question. The arc can be rotated about its diameter and the carrier about the arc so that any direct straight path to the lesion may be achieved.

The base of the frame contains pins, slots and recesses into which the base of the marker plates fit exactly so that their positions are accurately known in relation to the frame [72]. From knowledge of the positions of the reference points, the coordinates of structures in each image are measured accurately with respect to the frame coordinate system [72].

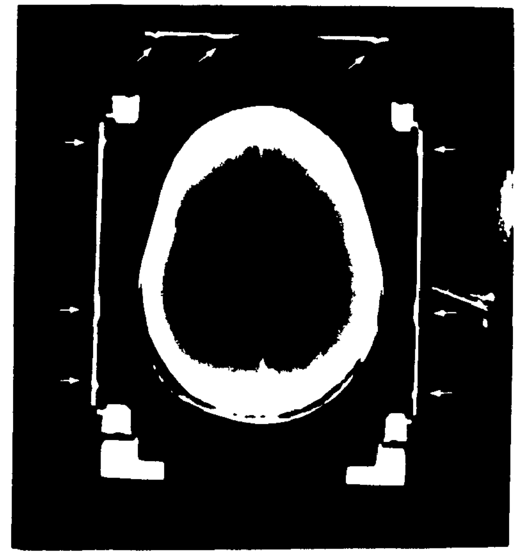

Figure 2.3: Typical axial MR (left) and CT (right) tomographic slice. Arrows indicate location of fiducial markers appearing in image.

2.3.2 Tomographic fiducial markers

To enable the coordinates in the tomographic images to be correctly identified, the scaling, position and orientation of the 2D slice within the 3D frame must be determined. If three non-collinear points are localized on the image with respect to the frame, the parametric equation of the image plane may be determined. This is achieved by the use of Z or N shaped fiducial markers embedded in plexiglass plates which are attached to the sides of the frame. A cross-sectional slice through the frame displays one set of three points for each marker plate intersected (see Fig. 2.3). Analysis of these points in the image, combined with the knowledge of the frame geometry, enables the location of the middle point (from the diagonal member) to be known with respect to the frame coordinate system. A cross section containing the intersection of three plates allows the location, scale and orientation of the slice to be determined exactly. A simple transformation matrix relates the 2D location of any pixel in the slice to its 3D frame coordinate [67,72].

The markers are made of different materials for each modality. The CT system employs *N*-shaped aluminum markers. Only three plates, one on each side and one in the front of the frame, are needed for transverse imaging in CT. The MRI system uses closed channels in the same *N*-shape filled with a solution of 0.7g/l of $CuSO_4$ (copper sulfate). Five plates are needed for the MRI system since sagittal and coronal images are also allowed.

2.3.3 DSA fiducial markers

Because of the different geometry, another type of marker must be used for the DSA system. Small steel pellets are embedded at the corners of a square on the four plates used, front and back for PA images, left and right for lateral imaging. Since the DSA technique yields projections of the brain rather than cross-sections, where eight marker points are seen in each of the PA or lateral view. Although these points are placed at the corners of similar squares on opposite sides of the frame, they appear in the image to be on trapezoids of different sizes (see Fig. 2.4). These points nevertheless precisely define the beam geometry of the system.

Calculation of the homogeneous transformation matrix [88] defining the projection image geometry is achieved from a minimum of six identified fiducial markers using a least-squares minimization technique [92]. A single pixel in the image lies on an infinite ray passing through the focal spot. If the structure can be identified in another view, a simple intersection between the two rays will yield the 3D location of the point in the frame coordinates.

Figure 2.4: Typical left lateral DSA radiographic projection. Arrows indicate location of fiducial markers appearing in image.

2.4 Stereoscopic workstation

The stereoscopic images are viewed on a Tektronix SGA[6] system which has a single high performance monitor operating at a refresh rate of 120 Hz. The monitor is equipped with a liquid crystal polarizing shutter, which is synchronized with the field rate, to display right- and left-eye images alternately while changing the circular polarization of the transmitted light between the left and right-handed sense (see Fig. 2.5). The operator views the stereo pair via correspondingly polarized glasses.

My colleague Chris Henri is working to extend the stereotactic environment described above to the stereoscopic workstation. At this time, cursors and vectors can be drawn interactively in stereo using a hand-held mouse and have been used to identify target points in 3D or simulate surgical probes and needles. Using stereo triangulation, the

[6]Tektronix Stereo Graphics Adapter.

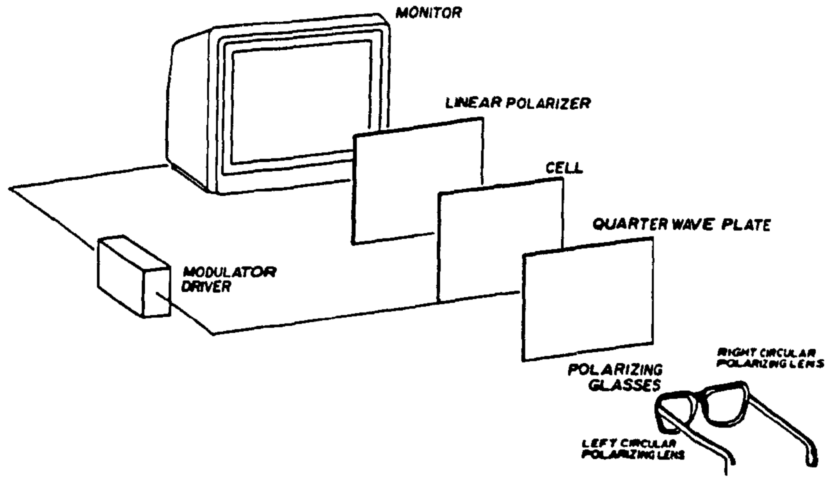

Figure 2.5: Exploded view of Tektronix stereo display system.

3D location of the cursor can be displayed to the user in real time. The accuracy with which the depth coordinate may be determined is a function of pixel size, stereoscopic image disparity, and the ability of the human observer to place a cursor accurately at the desired depth using the computer system and mouse. It will be necessary to continue the work described in [36] to establish limits on accuracy and precision.

2.5 Summary

Each imaging technique provides information complementary to the others. The goal of this work is to provide a basis for the integration of images from multiple modalities. The imaging geometry, defined by the location of the fiducial markers on the stereotactic frame in the DSA image, is used to create matched tomographic volume-rendered projections. This ultimately provides the neurosurgeon with 3D stereo images containing both vasculature and anatomy that can be used in the stereotactic planning software.

Chapter 3

Review of rendering methods

A volumetric data set is formed by sampling (i.e. measuring) a volume of interest at regular intervals in 3D space. Each volume element is called a ***voxel*** and is defined by the intersection of six planes, two adjacent planes from each of the three sets, and are typically arranged in a 3D lattice. Each voxel has an (x, y, z) position as well as one or more measured parameters associated with it. Tomographic scanning techniques yield a series of 2D images, with each one representing the sampled values in a 2D plane. A sequence of tomograms, stacked in order on top of each other, forms the volumetric data set. (see Fig. 3.1).

The procedure used to create a 3D model of an object from its representation on sliced images is called reconstruction. The goal of the 3D reconstruction algorithm is to find the 3D structure from the 2D information on each slice and display a 3D view of the model on a 2D screen.

Approaches to the display of 3D data have been:

1. serial presentation of 2D slice images from the 3D volume (multiplanar reformatting).

2. multiple stacked 2D images.

3. wireframe representation.

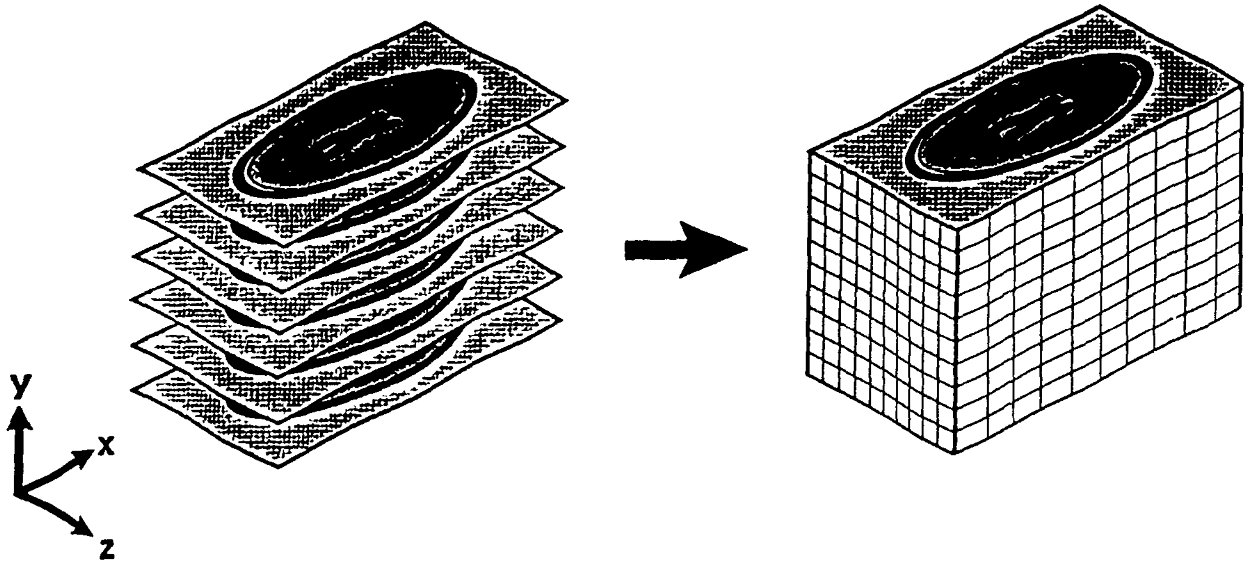

Figure 3.1: A volumetric data set is a 3D array of cells called voxels.

4. surface rendered views.

5. volume rendered views.

The following summary of these techniques begins with the simple surface rendering techniques, followed by a description of the cuberille model and finally the ray tracing and direct volume visualization methods.

There are three basic techniques for rendering 3D volumes; cross-sectional rendering (or multiplanar reformatting), threshold (or surface) rendering and non-thresholding direct volume visualization. In cross sectional rendering, all voxels are considered to be completely opaque. and the user may interactively remove sections of the data by slicing through it with an arbitrary cut-plane. As the cut-plane moves through the volume, the voxels falling on the slice are texture mapped onto the corresponding surface of the visualization cube. Programs of this kind are PIXAR's Cubetool[1] and the

[1] PIXAR inc., San Raphael California.

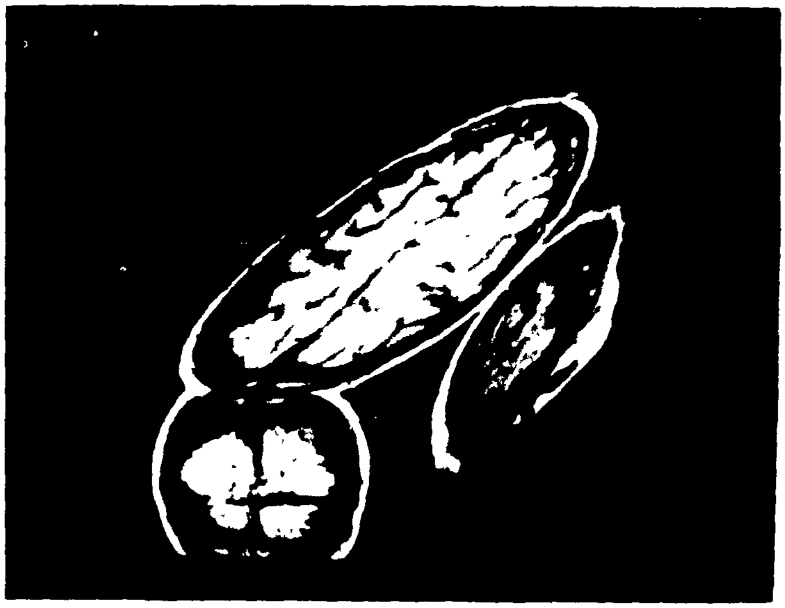

Figure 3.2: Example of an MR data set visualized with the Cubetool software. The grey values from the volume are mapped onto the surface of a 'cube' that is manipulated by the user.

TAAC's VOXVU[2] (see Fig. 3.2). In threshold rendering, all voxels falling outside the threshold region are considered to be completely transparent and do not contribute to the rendered image. The resulting image is a view of the surface corresponding to the threshold interface. There are a number of methods used to render these iso-surfaces[3]; surface reconstruction, ordered traversal (front-to-back or back-to-front) of the data and ray-tracing. These three schemes attempt to fit surfaces to the data and then render the surfaces visible using standard graphics techniques. Direct volume visualization techniques differ by rendering images directly from the volumetric data, projecting each voxel onto the image plane. The projection image is formed by compositing or merging the voxel values, one on top of each other.

[2]SUN inc., Transcept Application Accelerator.

[3]An iso-surface is the 3D equivalent of an iso-contour in a 2D image. All points on the surface have the same value.

3.1 Surface rendering

Surface rendering of volumetric data represents only the surface of structures. The first group of algorithms manipulates surface contours as the basic data structure. These contours may be created manually by outlining the structures of interest, or created automatically by thresholding and iso-contour following routines. Many techniques are available to represent 3D structure in a 2D image such as perspective, lighting, shading, hidden surface elimination and translucency.

The first scheme used in the past by Newman to view the reconstruction of individual contours begins by thresholding the slice data and extruding the resultant image in the axial direction by the slice-to-slice distance. These slabs are then stacked together and viewed using simple geometric routines, giving a staircase-like representation of the thresholded structure [63] (see Fig. 3.3).

3.1.1 Surface tesselation from iso-contours

One of the first methods used to represent iso-surfaces was to draw a vector mesh by hand, connecting iso-contours on consecutive slices of parallel planes within the data volume [103]. Algorithms were developed to automatically generate polygonal surfaces from these stacks of 2D contours (see Fig. 3.3). Most tesselation methods are based on triangulating the 3D surface ribbon between two consecutive contours, differing mainly in the choices made during triangulation.

Keppel was one of the first to devise a method to reconstruct the ribbon surface defined by two consecutive contours [44]. All possible triangular tiling arrangements are associated with a directed graph, and the graph is searched to identify a minimum cost path corresponding to a surface that maximizes the enclosed volume. If more than one contour exists on a slice, ambiguities arise when determining which contours to connect. Fuchs et al. presented a solution to the same problem that minimized the surface area of the tesselated ribbon [24]. However, their algorithm is only able to deal

Figure 3.3: On the left, contours defined by hand on an MR data set are shown stacked, one on top of the next. On the right, the contours have been tesselated and surface rendered.

with simply branching structures. Christiansen and Sederburg achieve a similar result by mapping the contours onto a unit square [12]. In this case, the triangulation is selected on the basis of the shortest diagonal and ambiguous branching is resolved by user interaction.

These tesselated surfaces are displayed on a vector graphics terminal using wireframe in perspective, with or without hidden lines to aid in depth perception. When displayed on a raster terminal, the polygons are displayed as opaque tiles with lighting highlights and shadows.

Suetens et al. [93] have developed a stereoscopic workstation where the objects, once defined, are displayed in stereo using a wireframe model. This system relies on motion parallax to give the impression of depth. As the user moves in front of the terminal, or as the user manipulates the object, its projection is recalculated, taking into consideration the viewer's new position, giving a pseudo-holographic impression.

Since there exists an infinite number of 2D surfaces passing through a given set of closed curves, Xu and Lu solve the surface reconstruction problem by thinking of it in terms of a functional minimization problem, which is solved by a partial differential equation method [104]. This is possible when the constraint is imposed that the sum of the curvatures of the surface should be at a minimum. This is equivalent to stating that the surface be maximally smooth.

In 1976, Mazziotta et al. developed the THREAD[4] system [61], capable of displaying 3D images from 2D sections by using threshold segmentation (with manual correction) and an automatic boundary detection scheme to outline the area on interest on the 2D images. The sectioned object outlines are viewed at different angles with and without hidden-line removal. Surfaces are defined and viewed in the same manner by tesselating the surface between the outlines on the consecutive slices and applying a simple surface-normal based shading algorithm to the surface polygons.

3.1.2 Surface patches

Surface patches is the term used to describe the class of surface representation using non-linear polynomials or splines to map a given surface. Surface patches are conceptually simple and render curved surfaces naturally. The most widely used mappings are the B-spline and the bi-cubic families. The tesselation methods previously described can also be used to define the control points (or nodes) of the surface patch. Catmull [10] has devised a sub-division procedure to display non-linear patches that recursively divide the region into sub-patches until the transformed sub-region covers only one pixel. Scan-line implementations of this algorithm have been developed by Blinn and Lane-Carpenter [48].

[4]THREAD stands for THree dimensional Reconstruction And Display.

3.2 Volumetric rendering techniques

Increased machine capacity for data storage and processing has made it possible to develop algorithms that examine the data in its natural (voxel) format instead of resorting to simplification and compression of the data using explicit surfaces. These techniques can be considered as either forward mapping or backward mapping schemes.

Backward mapping methods use computer graphic techniques of ray-tracing to form the image. For each point (or pixel) in the rendering, a ray is projected through the data volume, and each data element (voxel) intersecting the ray influences the value ascribed to the image pixel.

Forward mapping algorithms turn the problem around and project each voxel directly onto the image plane, relying on the traversal order of the data to ensure proper hidden surface removal.

A volumetric data set is intrinsically spatially presorted, and this property makes volume rendering feasible. For any viewing direction, one can easily determine the foremost corner voxel. The data traversal order can be established without explicitly sorting the data simply by stepping through the rows, columns and slices in the order that will ensure that the foremost voxel is visited last. (One only has to determine the traversal direction along each of the major axes.)

3.2.1 The cuberille model

Most volumetric algorithms manipulate a cuberille as their basic data structure as defined first by Herman et al. [38]. A *cuberille* is the subset of voxels in the data volume that belong to an object (or organ). The goal of volumetric rendering is to segment the volume into a cuberille representing the structure of interest and to project each of the segmented voxels to the image plane.

The representation of the surface of a medical object usually contains a large number of voxels (and consequently a large number of surfaces to be projected). Three attributes can be identified that accelerate the display process. The first of these is that each voxel has six sides, of which only three are visible from any one point of view [3]. The second, surfaces of the voxels on the inside of the structure are not visible, since they are hidden by the voxels on the surface of the structure. Thus, identifying and displaying only the boundary voxels will result in a faster display. Finally, the intrinsic ordering applies to the cuberille since it is a subset of the volume.

From the 3D array of numbers which constitute the series of image slices produced by a tomographic scanner, a particular organ can be selected using a 3D automatic boundary detection scheme. To extract a specific tissue of interest from an intricate data set, Lorenson et al. [56,13] use a connectivity algorithm. A surface detection algorithm described by Udupa [98] follows the volume elements of the cuberille on the surface of the object of interest, producing a large set of parallelepipeds whose faces describe an approximation to the true surface. The faces are displayed using a z-buffer algorithm[5] [63] with appropriate shading of the faces.

3.2.2 Shading and the cuberille model

Shading cuberille surfaces is not straightforward. The discrete nature of the model makes it difficult to estimate the normal to the true surface. From the comments above, there are only six possible orientations of the voxel faces. Naive application of a standard polyhedral shading model results in aliasing evidenced by banding in the shaded images. Herman et al. apply low pass filtering to the final images to reduce the effect of banding at the price of reducing resolution. Other solutions are offered in a paper by Chen et al. [11]. Context normal shading uses estimates of the surface normal based on voxel patches when displaying pre-determined surfaces, however, the most realistic images result when the gradient, calculated from neighboring cuberille

[5] The faces are sorted with respect to their z-coordinate and displayed in order so that the rearmost faces are displayed first.

elements, is used to shade a point on the surface of the model [31].

3.2.3 Front-to-Back rendering

The *front-to-back* algorithm merges three techniques to produce images quickly: front-to-back traversal of the data, use of the *dynamic screen* data structure and choice of rotational transformations so that a projected line segment of object is parallel to an image scanline [85]. The front-to-back traversal of the data implicitly eliminates hidden structures. The dynamic screen is a data structure that stores, for each scanline, the free (unused) runs of pixels that can make up an image. In this way, only voxels that project to the free area are manipulated. This technique produces images quickly since much of the data can be ignored using object and scan-line coherence. Gradient and depth shading is applied to the resulting image when the depth of each projected voxel is stored with the pixel in the image buffer.

A special multiprocessor architecture has been developed by Goldwasser et al. to implement the front-to-back algorithm [29]. It produces a shaded graphics display in real time with rotation, scaling, slice planes and shading transformations on grey-scale data.

3.2.4 Marching Cubes

The *marching cubes* algorithm [56], converts an array of data elements into polygonal format by the tesselation of a cube of eight adjacent voxels (four each from two consecutive slices). The data value from each voxel is placed at a vertex of the cube. If one examines the number of possible ways a surface can intersect a cube, only 15 are topologically distinct. The algorithm stores these tesselated configurations in a table that contains the edge information for the tesselation triangles. This logical cube is *marched* through the volumetric data set and each voxel is tesselated by retrieving edges from the pre-calculated table. The densities at the triangle vertices are tri-linearly interpolated from the original voxel data. The triangles (which lie on the surface of interest) are displayed with gradient and depth shading. Hidden surfaces are removed via scan conversion with a depth buffer, and polygonal surfaces are displayed using Gouraud shading. In addition to a single iso-surface being displayed, the algorithm allows any number of transparent surfaces, capped with triangles or surfaces mapped with the interpolated data from the original volume. The image quality is very high considering the complexity of the surfaces that can be rendered with this method. However the marching cubes algorithm is more costly in computing time than those previously described.

The *dividing cubes* algorithm was developed to eliminate the scan conversion step of the marching cubes algorithm [13]. As the number of the triangular facets increases, the size of each triangle decreases and eventually approaches the size of a single pixel. The dividing cube algorithm subdivides voxels that lie on the surface into small cubes, and then projects the intensity calculated for a triangle corresponding to the intersection of the small cube and the surface onto the viewing plane forming a gradient shaded representation of the object's 3D surface.

Adjacent voxels can be merged together in an octree, a structure for the representation of 3D objects that is an extension of the quadtree representation of 2D images[6]. This compresses the original voxel data and reduces the processing requirements for display [62] at the expense of pre-processing required to prepare and build the data structure. In addition, octree encoding achieves less data compression with grey-level data than with binary data. Also the traversal of the tree data structure involves overhead on a conventional computer [29].

3.2.5 Problems with surface fitting

Most of these techniques described have the following characteristics: The data are segmented into regions, geometric primitives are fit to the boundaries of the regions

[6]A quadtree is the 2D extension of a binary tree. An octree extends the concept to 3D.

and finally, classical techniques of computer graphics are used to display the primitives. These schemes differ mostly in the choice of graphical primitive, but suffer from the same problem, i.e. at some point in the visualization algorithm, a binary decision must be made. Either a surface exists at the voxel in question, or it does not. This decision leads to artifacts in the resulting rendered image.

Each of the surface extracting and display techniques suffers from the fact that much of the information in the original data is thrown away, that is one can no longer look inside the original volume since it has been reduced to a hollow polygonal shell. Each algorithm tries to retrieve some part of the lost information. For example, the connected contour routines attempt to recreate the connectivity existing in the original volumetric data; the cuberille approaches attempt to recover shading information from the binary object; and ray casting methods attempt to use depth shading alone or with an unnormalized estimate of the gradient.

3.2.6 Ray-tracing

Ray-tracing provides a direct method for identifying the visible points in a volume by simulating the effect of light rays. If the objects in the scene are opaque, then light cannot pass through them and will be stopped at the surface of the object. The visible points are those first hit by the ray traveling from the light source. In the simplest case, the rays travel from the source, are reflected off the object and impinge on to the image plane. Each pixel in the image has a corresponding ray (see fig. 3.4).

Binary objects – a voxel has a value of one if it belongs to the object and a value of zero if it does not – are displayed directly by method of Tuy et al. [97]. Their technique is different from those cuberille methods previously described, in that work progresses from the screen toward the object rather than the inverse. From each pixel in the image, a ray is cast through the data to the viewing point, and the voxels are tested to identify when the ray passes from the exterior to the interior of the object. Images are displayed in black and white and are shaded by depth, light direction and estimate of the surface

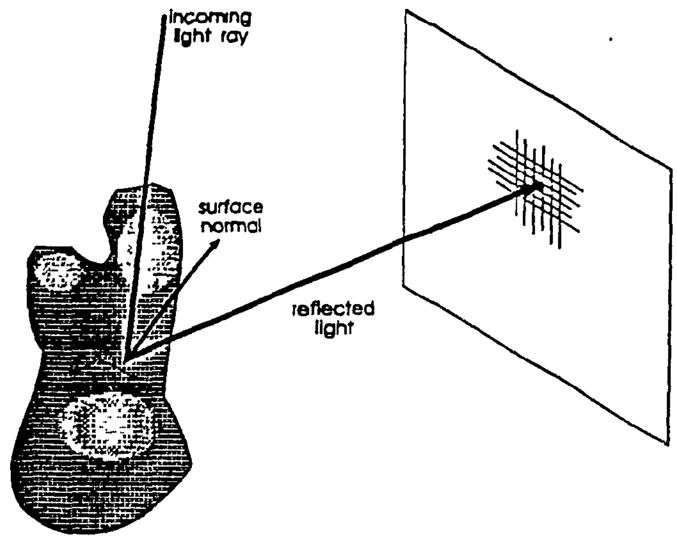

Figure 3.4: A light ray impinging on a pixel is modulated by the angle at which it is reflected from a surface.

normal. This algorithm has the advantages that interpolation and boundary detection do not have to be applied.

Farrel et al. stack the cross sectional images created by thresholding and display them using colour coding and depth shading by ray-casting [21]. However, the lack of surface shading and light reflection in the ray model makes it difficult to appreciate some structural details.

Hohne et al. have developed other methods that directly image the volume of data [40]. Their technique, *additive reprojection*, is a ray-tracing technique that computes the image by averaging the intensities of the voxels along rays through the volume onto the image plane, with the resulting image being equivalent to that of an X-ray projection through the volume of data. This is one of their *generalized projection* images.

This same group has used grey-scale data to produce surface shading based on the partial volume effects [39]. The relative volumes indicate the surface orientation, so that the grey level gradient can be used as a measure of inclination and thus produce the *projection of a surface*. *Cut plane projections* are produced by mapping the grey-level values from the plane of interest directly onto the image plane.

Additive projection techniques suffer from a confused or unclear appearance. Since all of the information present in the volume is projected onto the the image plane, there is be very little contrast available to distinguish gradual variations or small abrupt changes in the data.

Russel believes that *obscuration* plays an important role in visualization [89]. Pictures with complete obscuration (completely opaque) hide too much and images with little or none do not have enough contrast to be informative. They consider that selective use of partial obscuration is the key to producing informative images. This approach has been applied in a technique developed by Vannier et al. [100]. The *source-attenuation reprojection* method assigns a source strength and an attenuation coefficient (also known as *opacity*) to each voxel. Vannier has developed algorithms that accomplish depth shading by tracing rays through the volume array until they hit a surface, and then assign an intensity inversely proportional to the distance traveled along the ray. A surface is defined as a significant change in opacity value, so this approach is still threshold-based and suffers from the associated artifacts.

3.3 Direct volume visualization

In all of the algorithms previously mentioned, the important classification step determines what is seen in the final rendered image. The thresholding scheme defines the data left to be rendered, while the contouring method defines the outlines through which the surface model must pass. Any error in the classification is therefore carried through into the 3D model presented to the user. If the segmentation algorithm makes small mistakes in the estimation of the boundary of the object, then the approximation of the gradient of the surface will be incorrect. The surface normal will also be in error, and the lighting and shading model applied to the surface of the model will not result in an accurate representation of the true underlying structure.

This has led to the design of direct rendering algorithms where the classification and gradient estimation are done independently and in parallel [54,53,16]. For each point in the data set, the gradient is calculated from local information in small neighborhoods surrounding each voxel. In this fashion, the gradient does not depend on the classification scheme and is not affected by errors in the classification.

The term *direct volume visualization* was coined by Mark Levoy [53,54]. His algorithm is based on directly shading each data sample and projecting it on to the image plane. A shading model is applied to each data sample using local gradients to estimate surface normals. Continuous classification and enhancement operators are also applied to each voxel, but no explicit surface detection or fitting of geometric primitives is applied. The resulting colours and opacities are merged using volumetric compositing[7] from back to front along viewing rays. Since each stage in the direct volume visualization algorithm preserves the continuity of the data, non-linear operators such as thresholding are avoided.

The key advantage of this technique is that weak or ill-defined surfaces are not lost in a binary classification scheme. Since classification and shading are separate, the accuracy of the shading (and consequently the apparent orientation of the surfaces) does not depend on the classification. This allows the classification constraints to be more relaxed for volumetric techniques than when using geometric primitives. Another advantage is that the user performs the task of interpretation, while the processing simply enhances the features of interest. Rendering images directly from the volumetric data set eliminates the extensive preprocessing step necessary for surface based models, and all of the information in the original data is also preserved.

The technique used in this research is a modified version of the direct volume visualization rendering method developed at PIXAR (San Raphael, California [1]) and is

[7]Compositing is described in chapter 5.

described in detail in the next chapter.

3.4 Other rendering techniques

Other reconstruction techniques create real or virtual 3D images directly, rather than 2D projection on a screen. These methods include holography and the use of a varifocal mirror.

3.4.1 Holography

Holography is a form of lensless photography based on the recording of amplitude and phase of light illuminating an object [27]. To form a transmission hologram, a laser beam is used as a coherent light source and is split into a reference beam and an object beam. The object beam illuminates the object and is reflected onto the recording medium. The reference beam shines directly on the film. The two beams produce an interference pattern that is recorded on the film. When the developed hologram is illuminated by the reference beam at the same angle, an image of the original object is formed. The light waves traveling through the hologram are modulated by the recorded interference pattern and appear as if the viewer were looking at the original object.

Computed holograms are created by calculating the interference pattern and recording it on the film [42]. Unfortunately, output devices are not able to transfer the pattern to the film at the required resolution. Multiplex holograms have been developed to overcome this problem. In this technique, a series of perspective computer generated images are rendered and stored on the same hologram using multiple exposures[8].

Since they achieve a very high level of realism, holograms have been proposed as the solution for recording and viewing 3D objects [26]. Multiplanar holography can show semitransparency but it is difficult to represent colour within the volume and the

[8]Benton (87,88) uses approximately 1000 images in his Alcove holograms.

holograms suffer from insufficient luminous intensity. Fully computed holograms are not yet practical, since recording is difficult and expensive. Finally, the holograms may have to be viewed under special lighting conditions (point monochrome light source at a specific angle).

3.4.2 Varifocal mirror

The varifocal mirror eliminates the need for mental reconstruction by displaying the entire stack of tomographic images as a 3D image [33]. The display system uses a vibrating mirror with variable focal length optics. The mirror consists of an aluminized membrane stretched over a loud-speaker coupled to a CRT. The speaker is made to vibrate at 30 Hz and the tomographic images are displayed one at a time on the CRT in synchrony with the motion of the mirror. Because of the variable focal length of the mirror, the timing sequence of each frame and retinal integration, the stack of 2D images appear as a virtual 3D image in the mirror.

3.5 Summary and discussion

This chapter summarizes a number of the different rendering methods used to visualize 3D data sets. Most methods can be classified into either surface or volume rendering techniques.

The major advantage of surface representations is their simplicity, in that once created, they can be easily manipulated and visualized from any angle. Surface representations are appropriate when the goal of the visualization is to analyze an iso-surface corresponding to a specific threshold or the boundary layer between two regions, but it is not always useful when trying to understand the volume as a whole. An implicit assumption of these rendering algorithms is that the model of a thin envelope depicting a boundary or an iso-surface accurately represents the volume. Often, the interior of a structure contains mixtures of several different substances, and their local variations are lost if the volume is reduced to a series of layers. Reconstructed images of selected iso-surfaces are frequently difficult to interpret because many independent shells may be detected with the same threshold value. For example, both the cortical and epidermic external boundary can be extracted as surfaces using similar threshold values.

Producing a polyhedral model requires extensive preprocessing (often including operator intervention), and in this process, much of the information in the original data is lost. If there is a change in classification, or in the cut plane, the data must be re-processed.

Techniques exist to resolve the branching problem between slices when creating the surface model. The processes work well enough when the connection between contours is fairly obvious, but fail when branching becomes ambiguous, particularly when the inter-slice distance is much greater than the in-plane voxel dimensions. In this case one must resort to operator intervention [12] or heuristics [24].

The tesselation method is not always appropriate for diagnostic purposes. Only simple objects can be delineated and tesselated automatically. Objects whose boundaries can easily be identified using thresholds may be visualized using these techniques. However, the tesselation may not approximate the true exterior of the chosen structure, especially if the distance between slices is large. It is not reasonable to assume that the exact position and orientation of a small planar patch can be recovered from a point sample (such as a voxel) since this information is most probably lost in the sampling process. Hence, the resulting surface may not be an accurate depiction of the actual object.

Surface rendering usually yields a more recognizable, and hence a more comprehensible image than volume rendering techniques. This is an important consideration when the primary goal of the rendering is to communicate results to a wide audience[9].

Direct volumetric rendering techniques have been developed in an effort to resolve some of these problems. The most important improvement over the surface techniques

[9]Comment made by Craig Upson, Siggraph89

is that the binary classification step is eliminated. This allows poorly defined surfaces to appear (however slightly) in the rendered image.

The quality of the images depends critically on the inter-slice distance. When the slices are thick, and widely spaced, artifacts are readily seen and distract from the image. It is interesting to note that volumetric rendering techniques show only the information available in the original data. If the slices are spaced far apart, the polygonal tesselation display using Phong shading gives a result that is pleasing to the eye, only because it creates the illusion that there is more information in the data than there really is.

At this time, the state-of-the-art provides many techniques to create attractive images for qualitative evaluation. This thesis project takes one step in the direction of quantitative analysis by developing a technique that allows geometric measurements. By using stereoscopic display and a fiducial marker system during imaging, distance measurements can be made within the volume. The use of these fiducial landmarks also permits the rendering of volumetric data that can merged with DSA images, providing more information to the radiologist for analysis and diagnosis and to the surgeon for surgical planning.

Chapter 4

Rendering, Matching and Merging

Volumetric rendering techniques project the entire 3D volume onto the image plane as if each voxel were a semi-transparent gel filter. As described in detail in this chapter, the resulting image is strongly influenced by the characteristics assigned to each voxel.

The goal of this thesis is to develop a method, using these volumetric projection methods, to merge tomographic data with DSA images so that the resulting composition is sufficiently accurate to be used quantitatively in stereotactic neurosurgery planning.

This chapter begins with a short summary of the PIXAR rendering method, along with the modifications to the algorithm required to create semi-opaque direct volume renderings. Two types of merging are used to combine images from different modalities. The next section describes the first, a volumetric registration tool used to bring two or more volumes into correspondence, while the last section describes the second technique developed to precisely reproduce the geometry of the DSA projection in the volume rendering.

4.1 PIXAR rendering technique

The scheme used to render the volumetric medical data to create 2D projection images is based on the PIXAR volume rendering algorithm [17], and uses the PIXAR ChapVolume software library [2]. This library contains subroutines to manipulate rectangular arrays of pixels called ***pixel windows***. Facilities include geometric transformations (rotation, scaling, translations), mathematical operators (addition, subtraction, multiplication of pixel windows or constants), logical operators, lookup table manipulation and spatial filtering functions for 2D and 3D convolution.

The PIXAR method is a forward mapping, volumetric projection rendering technique, which creates an image by projecting each voxel onto the image plane. While the method permits the 2D visualization of structures in a 3D data set without explicit surface detection, a number of assumptions regarding the data must be satisfied to avoid aliasing in the projection image. The sample value at each voxel corresponds to the local measurement of a signal associated with some characteristic of the object, for example, tissue density in CT, proton density in MR or the density of a radioactive tracer in PET. It is assumed that the sampling frequency of this signal satisfies the Nyquist criteria [64]. If Nyquist sampling is not possible, it is assumed that the measurement process incorporates a low-pass filter before sampling the signal. The data set is also assumed to have isotropic resolution (i.e. made up of cubic voxels). This is not the case in medical imaging, where the inplane resolution is much higher than the slice thickness, violating these assumptions by creating oblong voxels. Interpolation techniques to deal with this fact are addressed later in the chapter.

This section begins by describing simple additive projections that average the voxel values along columns perpendicular to the image plane. The following sub-sections describe existing enhancements to project surface information and summarize segmentation techniques used by the algorithm. The section concludes with a discussion of the problems with the PIXAR model.

4.1.1 Simple projections

In general, a projection operator transforms points in a coordinate system of dimension n into points of a coordinate system of dimension less than n. In the case at hand, the

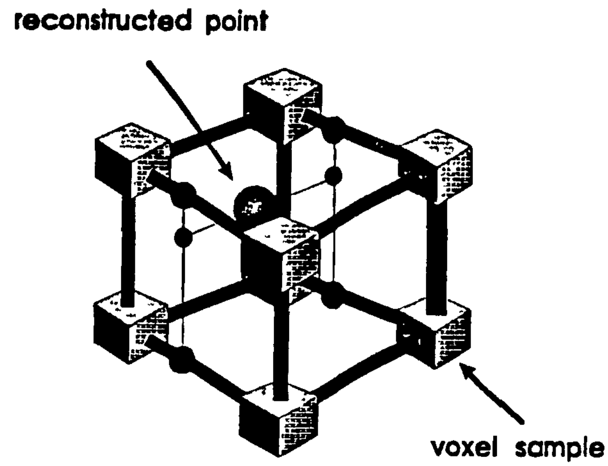

Figure 4.1: The value at some coordinate (x, y, z) is reconstructed by interpolating it value from the neighbouring samples.

projection is from the 3D stereotactic frame coordinate system present in the volumetric data set to 2D image plane coordinates. The PIXAR projection algorithm provides only for orthographic projections.

With a general ray-tracing technique, a ray is cast through the volume to a pixel on the image plane. All of the rays may be parallel (parallel or orthographic projection) or can emanate from a single point (the view-point or focal spot) to define a perspective projection. The view is defined by the direction of the rays or by the location of the viewpoint, and the projection is calculated by resampling the volume along these rays. These resampled points are reconstructed from neighbouring samples in the volume (see Fig. 4.1), a technique which requires arbitrary access to the data. While the architecture of the PIXAR does not allow for fast arbitrary access to the 3D volumetric samples, it is designed for quick projection along any of the three major axes. When projections along divergent rays are needed (to match DSA and projected tomographic data for example),

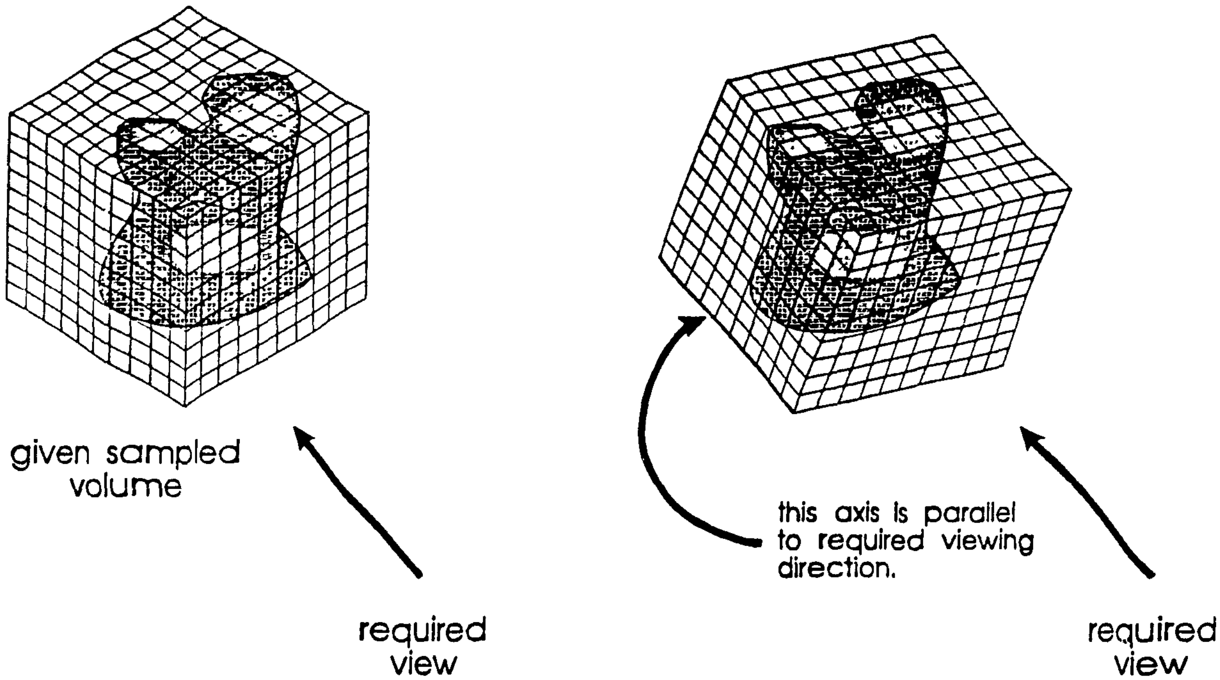

Figure 4.2: Pixar projection technique. On the left is a given sampled volume. The required view is indicated by the arrow. On the right is the resulting resampled volume. The indicated axis is parallel to the required viewing direction.

or when off-axis views are required, the data within the volume is manipulated so that orthographic projections, which only require re-sampling parallel to one of the volume's axes, form the same 2D image as would be obtained using a ray-tracing technique (see Fig. 4.2). These manipulations include rotations, translations and scaling.

The simplest renderings to create using the PIXAR algorithm are *additive projections.* These are formed by averaging the voxel values along columns perpendicular to the image plane. This type of image is rarely used, since it is difficult to see any detail within the image. Usually, only the silhouette of an object within the volume may be recognized.

The controlled use of transparency and opacity associated with each voxel determine the contrast in a projection image, which is the result of simulating the projection of light through a gel-like substance. Each voxel in the volumetric data set has a number of characteristics, the two most important for projection being: 1) The opacity of the gel filling the cubic voxel, measured by the parameter α, may block the transmission of light. 2) A source, placed within the voxel, emits coloured light along the direction of projection, with separate intensities in the red (R), green (G), and blue (B) channels.

If a voxel is completely transparent ($\alpha = 0$), like clear glass, it does not affect the light impinging on its back surface as the ray projects though it towards the image plane. A non-zero opacity will attenuate the light coming from behind by the fraction $1 - \alpha$.

The projection method uses a simple compositing scheme modeled after an optical film recorder [80]. This recursive technique is known as an *opacity-weighted integral.* Placing one voxel over another, the projection begins at the back of the volume and ends at the front. The OVER operator defines this operation on a voxel basis to composite a foreground voxel, A, with another directly behind it, B, giving the result C:

$$C = A \text{ OVER } B = \alpha_A A + (1 - \alpha_A)\alpha_B B. \quad (4.1)$$

The colour-opacity, (R, G, B, α), of C is used in place of the background voxel when the next foreground voxel is composited. This recursive operation continues until all voxels have been merged, from back to front. The result of the opacity-weighted integral, the final C, is then displayed as a coloured pixel, with an appropriate intensity, on the image plane.

For implementation, the α value resembles the forth element in a homogeneous coordinate system, and it is used to normalize the first three elements. Thus, a voxel whose colour is (R, G, B) and opacity (α) is represented by (r, g, b, α) where $r = R\alpha$, $b = B\alpha$ and $g = G\alpha$. Since all colours are stored pre-multiplied, two multiplication steps in the composition are saved.

In general, computing the path of a ray and calculating the intersections with the voxels is an expensive operation. For an orthogonal projection parallel to a volume axis, the projection simplifies to a multiplane merge of the slices parallel to the image plane. Thus, the orthogonal projection through the i^{th} plane may be expressed as eq. 4.2 from [80],

$$I_i = C_i \text{ OVER } I_{i+1}, \tag{4.2}$$

where I is the accumulated intensity, C_i is the colour-opacity of plane i and the final image is left in I_0.

Hidden surface removal is achieved by the traversal order of the data. Since orthographic projections are used, the traversal order must simply ensure that the columns of data (one column per pixel) are traversed from back to front for a z-buffer type of hidden surface rendering to be accomplished. (Opaque voxels in the front are overlayed on those behind, implicitly hiding them.)

When the data volume is rendered as described above, using only the opacity of each voxel, an *opacity-only* image is formed , a monochromatic density image similar to a digital radiograph. If the colour of each voxel is used in the projection, then a *colour-opacity* image is made, where structures may be differentiated by their colours. *Refractive index only* images, which resemble the projection of the surfaces of the different structures in the data set, are created by projecting the gradient of the voxel opacities. A shading model may be applied in the projection of the opacity gradients to create a *shaded refractive index* image, while adding the colour information to the gradient projections results in *coloured refractive index* images which are coloured surface images. A *shaded coloured refractive index* image is rendered by using the colour information modulated by gradient with a shading model in the projection pipeline.

4.1.2 Opacity-only projections

Opacity only projections of scalar data are the simplest to render, the opacity-weighted integral projection being applied to each of columns in the data volume. Voxels with significant opacity are brighter and obscure more of what is behind them, and the resulting pixel is assigned the value of the last opacity calculation. If colour is applied in the projection routine, then the colour of the pixel will be that resulting from the compositing, divided by the value of the resulting opacity.

Colour is useful for distinguishing the different tissues found in the volume, and colour-opacity images may be quickly rendered for noisy or low-resolution data sets where the surface information is difficult to extract. These images may also be used in the planning of the more complex time-consuming renderings since they can be completed in a matter of seconds. As described previously, opacity-only projections suffer from the same problem such as a muddied appearance due to lack of contrast as additive projections.

4.1.3 Surface projections

A surface internal to the volume may be defined as the interface between two regions of different density. For surface-projection rendering, one assumes that the visualization of these interfaces will yield enough information about the volume.

The output magnitude of a gradient operator, after application to the classified data volume, yields large values at the interfaces between constant regions and null values elsewhere. The change in density may be visualized by projecting through the matrix of these gradient magnitude values. The estimate of an interface between regions is defined at each voxel by a surface normal, $\vec{N}$, and a surface strength (or magnitude), $|\vec{N}|$, and is derived from the voxel density. The surface normal $\vec{N} = (N_x, N_y, N_z)$, and is approximated by a finite difference of neighbouring densities and the surface strength by $|\vec{N}|$, the L_2-norm of $\vec{N}$.

A refractive index-only image is obtained by projecting directly through the surface magnitude information. Since surfaces are defined by fractional values at each voxel, diffuse transitions that would not be possible with a binary value representation may be visualized. Surfaces may also appear to lie (or to be located) at sub-voxel accuracy.

Coloured refractive index images may be rendered by using the surface magnitude

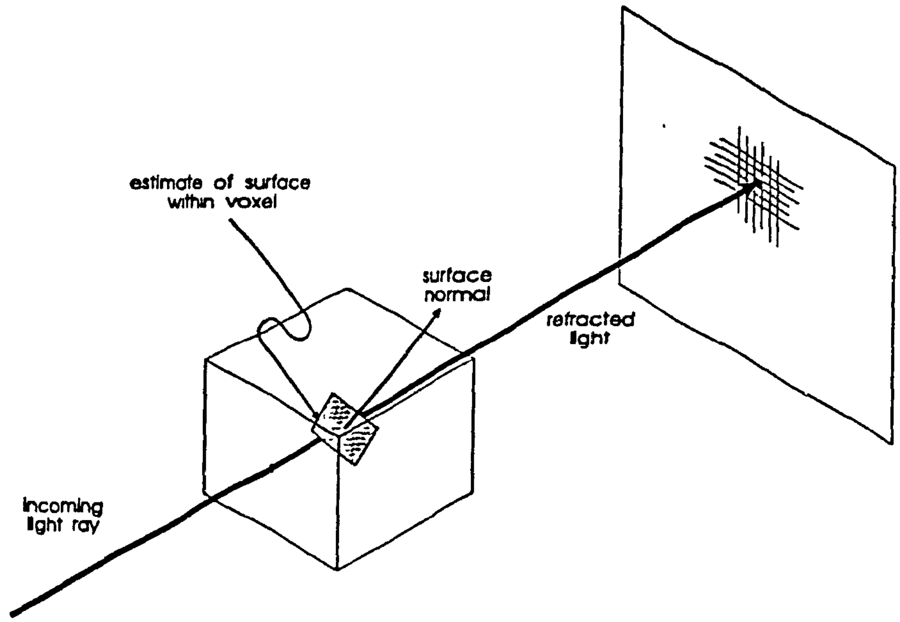

Figure 4.3: The light refracted through the voxel is modulated by the colour of the voxel as well as the magnitude and direction of the estimated surface normal.

information multiplied by the colour classified volume. When projected onto the image plane, the surfaces of the different materials may be differentiated by colour.

4.1.4 Shaded surface projections

In order to enhance the perception of depth in the rendered image, the 3D lighting cues may by applied to the data set. The finite difference given by $\vec{N}$ is used as the estimate of the surface normal direction for the lighting model.

The light enters the voxel from the back with an intensity I and an outgoing intensity I'. The intensity and colour of the light may be modulated by three voxel dependent effects, namely translucency, luminosity and refraction (see Fig 4.3). The direction of the surface normal may be used to modulate the intensity of the transmitted light refracted to the next voxel.

4.1.5 Classification

Visualization of 2D and 3D data shows structures known to exist in the brain. One is often more interested in a single, or a few, cerebral organs rather than the head as a whole. The purpose of segmentation is to partition an image space into meaningful regions [30], which should correspond to the different materials and structures of interest in the volume of data and the particular method should be able to deal with voxels that contain information from more than one type of tissue. This partial volume effect results in the value associated with the voxel (its density) being a weighted average of the substances within it. It is important therefore that a non-binary continuous classification scheme be used in order to avoid aliasing and contouring artifacts since the voxel does not belong to one group or the other. This way, subtle surfaces that occur at the interface between two similar materials are not lost.

Look-up tables

In the simplest form, a classification algorithm uses a single variable (i.e. the pixel/voxel intensity) as the input to the decision function. Using a nearest neighbor decision function [18], each pixel is classified as belonging to the set with the nearest center. This technique is easily implemented using a linear look-up table, where each of the input values is mapped through the pre-calculated table into material percentage values that correspond to the proportion of each material in the voxel. The look-up table may be based on the analysis of the image histogram, using its peaks as the centers of the density regions. This technique works well if the following of conditions hold:

1. Each material corresponds to a known density value.

2. Each voxel contains no more than two materials.

3. If two materials are found in a voxel, then their densities must follow each other in order in the linear table.

This classification is appropriate for neurological CT data. However, there are some voxels that do not adhere to the third criteria. When classifying voxels corresponding to the nasal passages, the air-bone interface violates the ordering of the table since soft-tissue is classified between air and bone. Also, most internal soft-tissue organs do not meet the adjacency criteria so this classification technique is not easily used for visualization without resorting to manual intervention to carve away unwanted tissues. In these cases, thin layers of the mistaken tissue appear at the interface between the two original substances.

Use of several variables

In order to improve the strength of discrepancy in the decision function, other image parameters may be used along with the pixel density in the decision function. The density of the T1 and T2 weighted magnetic resonance images may be used as the two decision variables. Classification now becomes a problem of locating clusters in a two-dimensional space. A scatter-plot (a histogram of two inputs) is used to locate the centers of the regions of interest. Once again, nearest neighbor classification are used where the distance function is defined as the length of the chord joining the center of each region with the 2D point corresponding to the voxel's parameters.

Classification is a difficult problem and although routines exist that have specific success, a generalized technique has yet to be developed. However, direct volume rendering techniques are not as sensitive to errors in classification as the polygonal surface rendering methods. The following sub-section describes how the output of the continuous classification scheme is used to segment structures of interest.

4.1.6 Matting

As was previously mentioned, obscuration (or the lack of it) plays an important role in visualization [S9]. By lessening the presence of certain regions of data, or by removing them altogether, other sub-volumes may become visible in the fore-ground of the rendered image. A *matte* is a mask volume, used to eliminate or enhance certain regions of the data volume.

A volume, V, may be combined with a matte, M, with the following operations [80],

$$V \text{ IN } M = MV, \tag{4.3}$$

$$V \text{ OUT } M = (1 - M)V \,. \tag{4.4}$$

At each voxel, the matting operation multiplies the value in the data by the value in the matte. The IN operator is a multiplicative intersection operation yielding the portion of V inside the matte M. The OUT operator yields the part of V outside the matte by complementing the matte M and then forming the intersection with V. Each value in the matte volume is a scalar fraction, defining the percentage of the matte in the voxel. These values are used to lower or raise the percentage of a material in the matted region or to change other properties of the material. By making the values fractional instead of binary, the boundaries of the defined regions remain smooth and continuous preserving the continuity of the data and avoiding potential artifacts in the renderings.

Enhancing data in the foreground and fading it in the background by depth shading is achieved by multiplying the original data volume by a ramped matte volume, thus making voxels near the viewer brighter than those farther away. In this way, *matte* volumes are used as operands with the original data in spatial set operations that modify the data.

The matte volumes may be created manually, for example, outlining the cortical rim in order to remove the scalp from the rendering, or geometrically, in order to define a wedge or cut-plane through the data to move or remove a sub-volume of data. The matte volumes may also be created from the output of the classification technique in order to segment structures identified by the classification algorithm.

4.2 Modifications to the PIXAR algorithm

The highly transparent images formed using the gel-model may not be appropriate for neurosurgical planning. The model of light-emitting particles, aligned on surface boundaries within a semi-transparent coloured gelatin, is not particularly intuitive. While the rendered images share some characteristics with radiographic images (back-lit, light passing through the object, high level of transparency, problems with contrast), they do not model the organs as they are seen during surgery. The human brain is opaque, and during surgery, it is lit by direct and ambient light, usually from the same side as the surgeon. If the visualization is treated as an abstract image, adherence to a common, familiar, underlying physical model is not necessary. However, the human visual system picks up on cues that resemble those from everyday experience. Inconsistent cues may give rise to ambiguous and sometimes unplausible interpretations. For these reason's, a number of modifications are introduced to the algorithm and are described in this section.

The surface normals estimated in the PIXAR technique are calculated from the voxel values in the classified volume. Although the technique is not as sensitive as surface rendering algorithms, errors in classification lead to local errors in the surface estimate. The modified technique computes the classification and the gradient volumes independently, similar to the rendering algorithm of Levoy [54]. The gradient is estimated by the finite difference in density of neighbouring voxels within the original data.

In order to enhance the perception of depth in the rendered image, the 3D cue of lighting due to refraction is applied to the data set in the PIXAR model. Only the component of the refracted light, parallel to the projection direction, is used in the opacity-weighted integral. A more realistic cue is given by surface orientation, perceived by the reflection of light from the object due to an externally positioned light source.

Polyhedral surface rendering techniques are well known in computer graphics [23,63,88]. The basic premise of these methods is that an object, represented by a polygonal model,

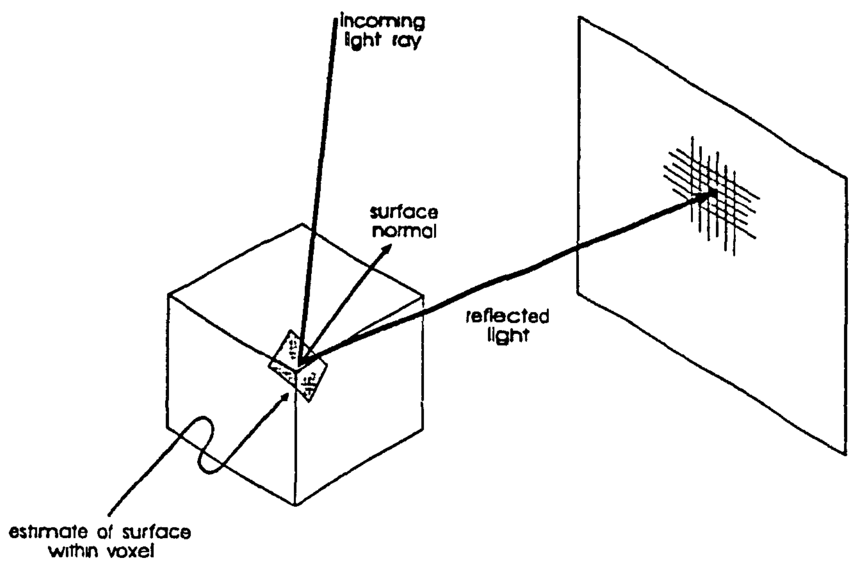

Figure 4.4: The intensity of the light reflected off the voxel is modulated by the cosine of the angle between the estimated surface normal and the direction of the incoming light as well as the normal's magnitude.

is lit by both ambient and direct light. Ambient light illuminates the object evenly in all directions, while the direct light is considered to eminate from a point source located at a specific position. The cosine of the the angle between the direction of the light rays and the normal to the polygonal surface modulate the intensity of the light reflected back to the viewer via the image plane. The goal of the modifications is to incorporate this lighting model in the direct volume rendering algorithm (see Fig. 4.4).

Computer-graphics rendering techniques require some polyhedral model of the scene to be viewed, i.e. the shape, size, orientation and position of the surfaces must be specified. Here, the polygons are the result of the surface normal estimate at each voxel using the finite difference between voxel values. The position and orientation of the polygon is given by the voxel location and the normal estimate and the size of the polygon determines the amount of light reflected back to the viewer and is estimated by the magnitude of the surface normal. The shape of the polygon is irrelevant, since the surface within each voxel is orthographically projected to a single pixel using the PIXAR software libraries.

The model yields fairly realistic, familiar images of objects, which appear as one is used to viewing them in everyday experience. Realism can be increased, at a cost of computational complexity, by adding shadows and reflections to the lighting model.

The existing PIXAR rendering algorithm assigns the same range of opacities to interior and surface voxels. Features that appear to lie on the rendered surface may actually be just above or just below it. Assigning a higher opacity to the surface voxels eliminates this error at the risk of hiding sub-surface structures that one wishes to visualize. The use of stereo projection images achieved by selecting rendering geometries that correspond to two stereo views with only a slightly higher opacity reduces both problems, since additional 3D cues are used to determine spatial relationships.

The algorithm is summarized in Fig. 4.5. Once the volume has been preprocessed to fit in the cuberille model (skew correction, interpolation and spatial filtering), the data within the volume is then oriented using translation, rotation and scaling transformations to form the required view. The projection routine begins by classifying the data in order to assign a colour and opacity to each voxel. Depending on the type of image required, colour is used to differentiate between the different classified materials in the scene. In parallel, the gradient is calculated at each sample point to estimate the strength and direction of the surface passing through the voxel. After classification, *mattes* are used to cut out, attenuate or emphasize certain regions of the data volume, and in particular, the gradient volume is used to accentuate the voxels with a large surface estimate and also to increase their opacity. The intensity of each voxel is modulated by the cosine of the angle between the surface normal and the lighting direction, which has the effect of making surfaces perpendicular to the light source brighter. Depth shading is achieved by multiplying the volume with a ramp, decreasing the intensity of voxels farther away from the viewer. The volume is then projected to the image plane

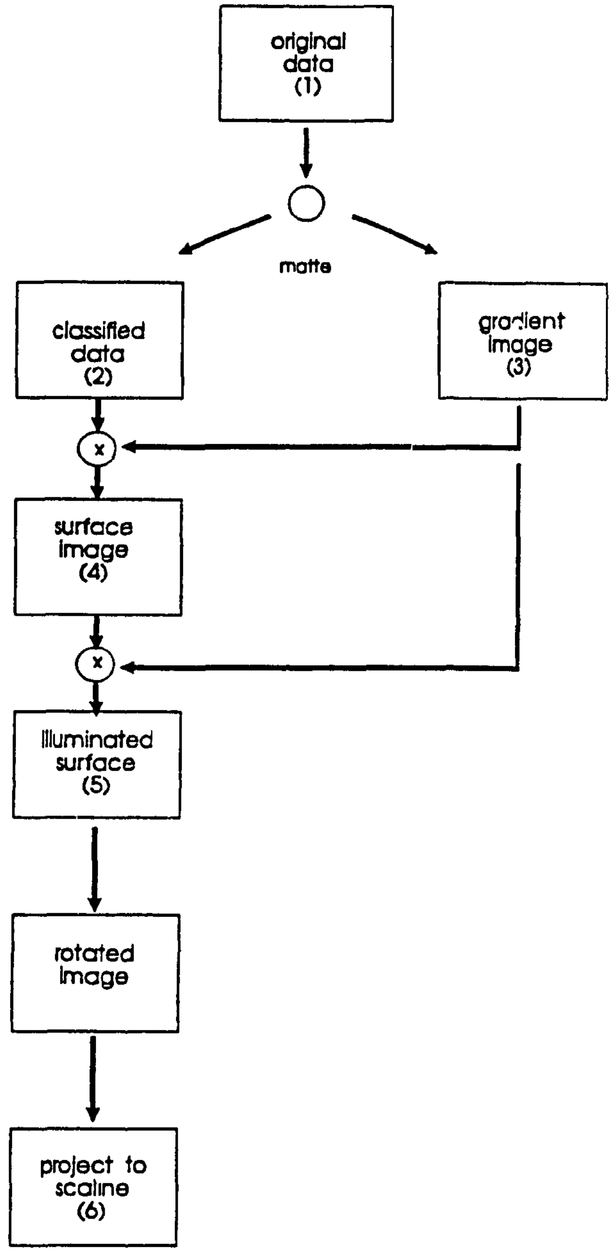

Figure 4.5: Flow chart of the rendering method. To create a single horizontal scan line n:

1. slice n is loaded and preprocessed (filtering, masking).
2. data is mapped through a continuous classification LUT giving **C**.
3. 3D gradient operator is applied to slices $n-1$ to $n+1$ giving **N**.
4. surface enhancement is achieved by multiplying **C** by $|\mathbf{N}|$ giving **S**.
5. lighting is applied by modulating **S** by the cosine of the angle between the lighting direction and $\vec{\mathbf{N}}$.
6. opacity weighted integration is used to project the slice to a single scan line.

using the recursive volumetric compositing rendering method.

4.3 Matching and registration

It is well established that CT, MR, PET (positron emission tomography) and DSA provide complementary diagnostic information. PET displays metabolic and physiological activity but has poor resolution and does not delineate anatomy. DSA images are non-tomographic but demonstrate fine details of the vasculature. MR and CT both image anatomy at high resolution but show complementary morphologies, e.g. calcifications and cortical bone are well seen on CT images while soft tissue structures are most clearly seen on MR images. It is often desirable to combine the complementary imaging information acquired during a complete study using the different modalities. When merging images, precise geometric correction must be applied to one data set so that the corresponding anatomical or morphological features in both images are assigned to the same spatial location.

Quantification of different parameters associated with specific regions of the brain is an important step in diagnosis. At the Montreal Neurological Institute, a region of interest brain atlas is used to classify structures in the brain [20]. Each region of the atlas is tailored by hand based on tomographic MR images. This atlas is then transformed to match the corresponding positron emission image data for quantitative analysis. This method is of key importance for reliable analysis of PET data and relies on the registration of the PET and MR tomographic volumes. Another application of image matching is in disease monitoring and grading for the the tracking of a lesion to detect changes of shape and size occurring over time for tumour growth or the change in average MS plaque size. This requires registration of multiple volumes of the same modality acquired at different times. A major difficulty in stereotactic neurosurgical planning is the determination of a vascular-free path to a tumour or lesion. The work accomplished for this thesis permits projection angiograms to be merged with tomographic data [14].

Before any merging or matching may begin, the volumetric data sets must be geometrically normalized by a number of preprocessing steps, described in the first part of this section. Two type of merging are then considered. The first is the registration of two or more volumetric data sets, while the second is the integration of a processed volumetric data set with a projection image. Different techniques are used by each approach.

4.3.1 Preprocessing

The assumptions stated at the beginning of this chapter apply to volumetric data sets that must be registered. If the volumes do not satisfy the requirements, some preprocessing of the data set must be completed to minimize errors in the resulting image before the matching, registration or rendering processes begin. The volume should be sampled isotropically and consist of cubic voxels. The preprocessing consists of an image by image correction of any distortions or artifacts, correction of any skew in the data, interpolation (resampling) to obtain cubic voxels and finally 3D spatial filtering to remove aliasing of high frequencies (if the sampling of a non-bandlimited signal is not proceeded by low-pass filtering).

Correction of skew

The volumetric data set is created by stacking the parallel 2D tomographic images as they are acquired, assuming the scan plane is perpendicular to the scanner axis. If the images are acquired obliquely, stacking them results in a vertical skew of the volume, and voxels that should be directly adjacent have a small offset between them (see Fig. 4.6).

This problem is corrected by introducing a relative translation within each slice by an amount determined by the gantry angle and the slice thickness, which may be calculated from the positions of the fiducial markers appearing in the tomographic images. Since a

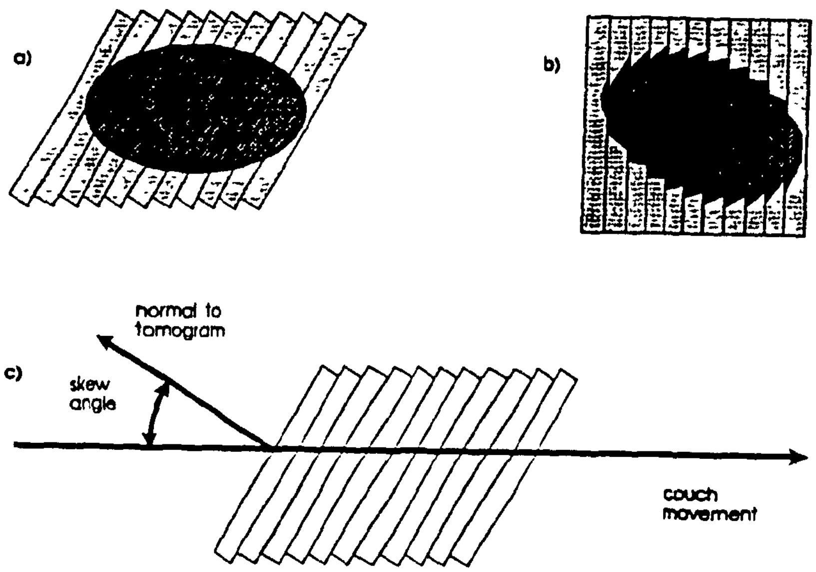

Figure 4.6: Volumetric Skew. **a)** shows slices acquired with angulated gantry. **b)** shows skewed volume when slices are simply stacked. **c)** angles defining correction amount.

particular pixel in every slice of the raw data corresponds to the same physical point in the scanner's field of view. The 3D stereotactic frame coordinates of this point in each slice defines a vector in the direction of scanner's couch movement. The fiducial markers in an image also allow a vector, normal to all slices, to be estimated in stereotactic frame coordinate system. The angle between the vector normal and the direction of movement is the angle of skew, and the projection of this angle onto the tomographic slice defines the amount of skew in the inplane directions to be corrected.

Interpolation

Medical tomographic data sets often fail to meet the requirements of isotropic sampling, since imaging constraints cause the slice thickness to be two or more times greater than the pixel size. In order to obtain cubic voxels, resampling of the original non-isotropic data set is achieved using tri-cubic interpolation.

4.3.2 Registration of volumetric data

Registration is the term used to describe the process of geometrically aligning two or more data sets so that the voxels in a specific area may be superimposed or merged. Two types of registration are generally considered: rigid alignment of the images and elastic registration of image sub-areas. Elastic registration distorts an image, warping it to match specific reference points between the two images, e.g. Ratib et al. [83] and Bookstien et al. use a model warping technique based on thin metal plate physics [5]. Since the data used in this study is relatively free from imaging distortions, elastic warping is not required.

Any technique for image correlation must take into account the different patient positions, the angle of scan with respect to the patient axis as well as the inplane and inter-plane resolution. Imaging requirements of the different modalities may require the use of different slice thicknesses, inter-section gap (if any), matrix size and pixel size. It is not uncommon to encounter transverse scans of the same patient performed on the same machine that differ by as much as 25 degrees in angulation. In this case, registration of the data sets cannot be achieved on an image by image basis and must therefore be completed in a 3D environment, where one data set must be re-sampled along the axes of the other. In order to correlate the two studies, a coordinate transformation must be found incorporating translation, scaling and rotation that maps corresponding points in each data set onto the same voxel location.

Three reference points are necessary to fix a rigid body in space, specifying the six degrees of freedom: translation along x-,y- and z-axes and three angular parameters of pitch, roll and yaw. Since the two volumetric data sets to be matched may also differ in scale, a fourth reference point is necessary. Registration techniques differ in their choice of measuring the distance between the points of reference, the calculation of the transformation and the application of the transformation to the data sets.

A number of different techniques have been used in the medical field to register data volumes. Pelizzari et al. have used an edge detection scheme along with surface fitting methods to register volumes of different modalities by what has been called the *head and hat* method [69]. The contours from the higher resolution data set, usually MR or CT, are used to define a *head* surface, while the *hat* points from the lower resolution data set, PET in this case, are fitted to the head by an algorithm that minimizes the mean distance between the hat points and the head surface. The residual r.m.s. distances between the points and the surface is on the order of 1 to 2 mm. This technique has found wide success [51,41,99] because it is conceptually simple and requires little expertise in neuroanatomy.

A volumetric registration tool has been developed in the NeuroImaging Lab at the MNI that incorporates interactive manipulation of the volumes with the Procrustes matching method [19]. In this program, the user specifies the transformation values for scaling, rotation and translation along any axis that are to be applied independently to each volume. All user input is accomplished with a hand-held mouse and the results are displayed to the user interactively. The three cardinal slices from each volume are displayed along with the resulting merged image by opacity weighted compositing. The user may pan through the data along either the cardinal axis or along arbitrary oblique directions.

The Procrustes matching algorithm has been partially implemented within the tool, which is used to identify corresponding points within the two volumes. The least-squares optimization is performed to calculate the affine transformation which minimizes the r.m.s. distance between the two groups of paired points. The identification of fifteen point pairs, matched with a standard deviation 3D uncertainty of 10 mm, resulted in a residual r.m.s. error of 1 mm between the merged volumes [19].

4.3.3 Integration of projection images with volume data

No single modality is suitable for imaging both anatomy and vasculature. Until new scanning techniques, such as MR angiography [79], become routine, the integration and correlation of vascular structures and anatomy will depend on the merging of DSA images with volume rendered projections of MR or CT data sets.

As described previously, the fiducial markers in the DSA images completely define the homogeneous perspective projection transformation matrix by the use of an optimization algorithm that minimizes in the least-square sense the difference between the projection of the fiducial markers in the image and those projected through the matrix [94]. This matrix fully describes the geometry of the image and may be used to determine the camera parameters, which are used in turn to project the volumetric data set onto an image plane that may be merged the the DSA image (see Fig 4.7).

Using a naive algorithm, the 3D position each voxel may be sent through the matrix, and its value mapped onto the projection location on the image plane. This approach suffers from the need to have arbitrary access to the data and aliasing would result in the projected image due to partial volume/pixel effects. Finally, hidden objects may

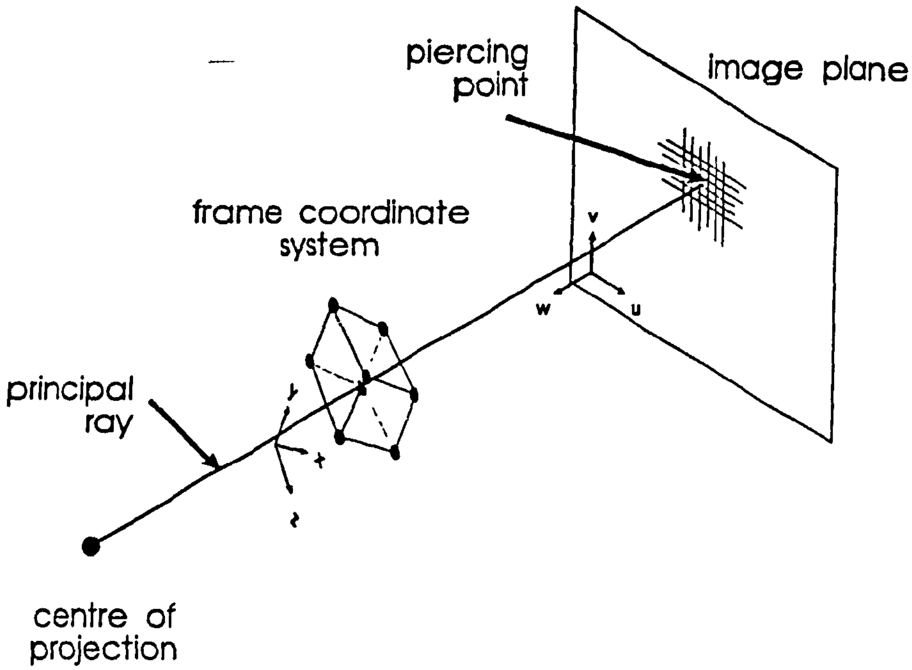

Figure 4.7: Projective imaging geometry.

not be properly obscurred depending on the order of traversal of the data. A direct volume rendering technique must be used.

A difficulty for any forward or backward mapping algorithm is perspective. The sampling rate of the volumetric data set with respect to the screen changes with depth. In order to make optimal use of the PIXAR imaging computer and maintain the simplicity of the orthographic projection scheme for arbitrary views, the viewing coordinate system remains fixed and the volume is oriented by geometrical transformations and resampled to lie in the fixed coordinate system.

Perspective views of the data are formed by creating a *truncated perspective viewing pyramid* as described in the next section. Once the volume has been resampled for the required viewing direction, the slices perpendicular to the viewing direction are resized according to the differential magnification required by the perspective (see Fig. 4.8).

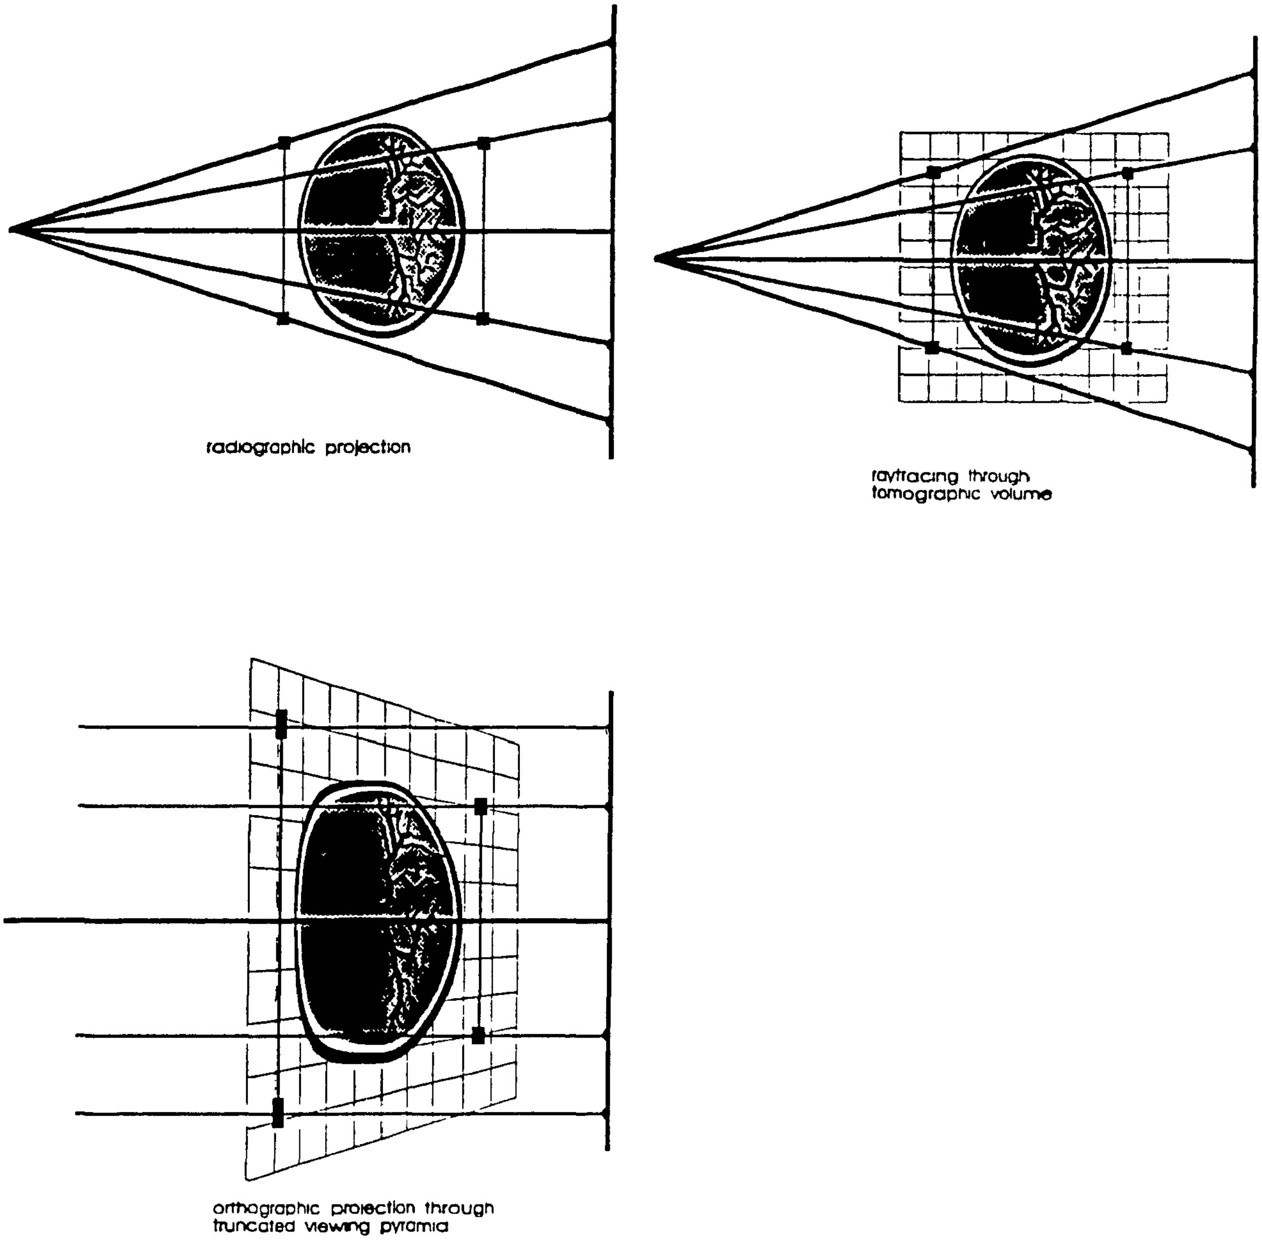

Figure 4.8: Truncated perspective viewing pyramid. **a)** shows the radiographic projection. **b)** shows how the same projection is achieved using ray-tracing. **c)** shows how the volume is differentially scaled to create a truncated viewing pyramid to achieve the same view using orthographic projections.

Creation of a particular perspective view:

Once the tomographic data set has been properly un-skewed and interpolation completed, the stereotactic frame may be at any arbitrary orientation within the working-volume. The DSA projection parameters must then be identified and the volumetric data re-oriented to match the projection geometry. The process is described below in a step-by-step fashion (see Fig. 4.9 for an overview).

1. The stereotactic fiducial markers are located on slices at the top and bottom of the volume and the location of the DSA markers are calculated (in voxel space). These DSA points are called the **A** group, and they define the orientation and position of the frame in the volumetric data set. It is necessary to re-position and re-orient the frame within the volume since orthogonal projections are used to yield the required view.

2. The homogeneous transformation matrix is decomposed into its component matrices using a technique due to Strat [92]. (Another technique is described by Ganapathy in [28].) Since the transformation matrix $\mathbf{M}$ is made up of translation $\mathbf{T}$, rotation $\mathbf{R}$, scaling $\mathbf{S}$, perspective projection $\mathbf{P}$ and cropping $\mathbf{C}$, it can be expressed as $\mathbf{M} = \mathbf{TRSPC}$.

3. The required orientation and position of the stereotactic frame are identified. Using the known locations of the fiducial markers in frame coordinates and the scanning parameters (slice thickness, slice spacing, pixel size), positions of the DSA markers are calculated so that the frame coordinate system is centered in the volume space and oriented to have its axes parallel with those of the volume. These points are called group B and are transformed by the $\mathbf{R}$ and $\mathbf{T}$ component matrices of the homogeneous transformation matrix calculated in Step 2. This transformed group of points is called the $\mathbf{B}'$ group and this defines the target position of the frame within the volume.

4. The Procrustes matching algorithm is used to identify the rotation and translation transformations necessary to bring the points of group **A** into alignment with the group $\mathbf{B}'$. (Scaling transformations are not necessary since the points of group **B** were created using the scanning parameters.) At this point in the algorithm, the volumetric data

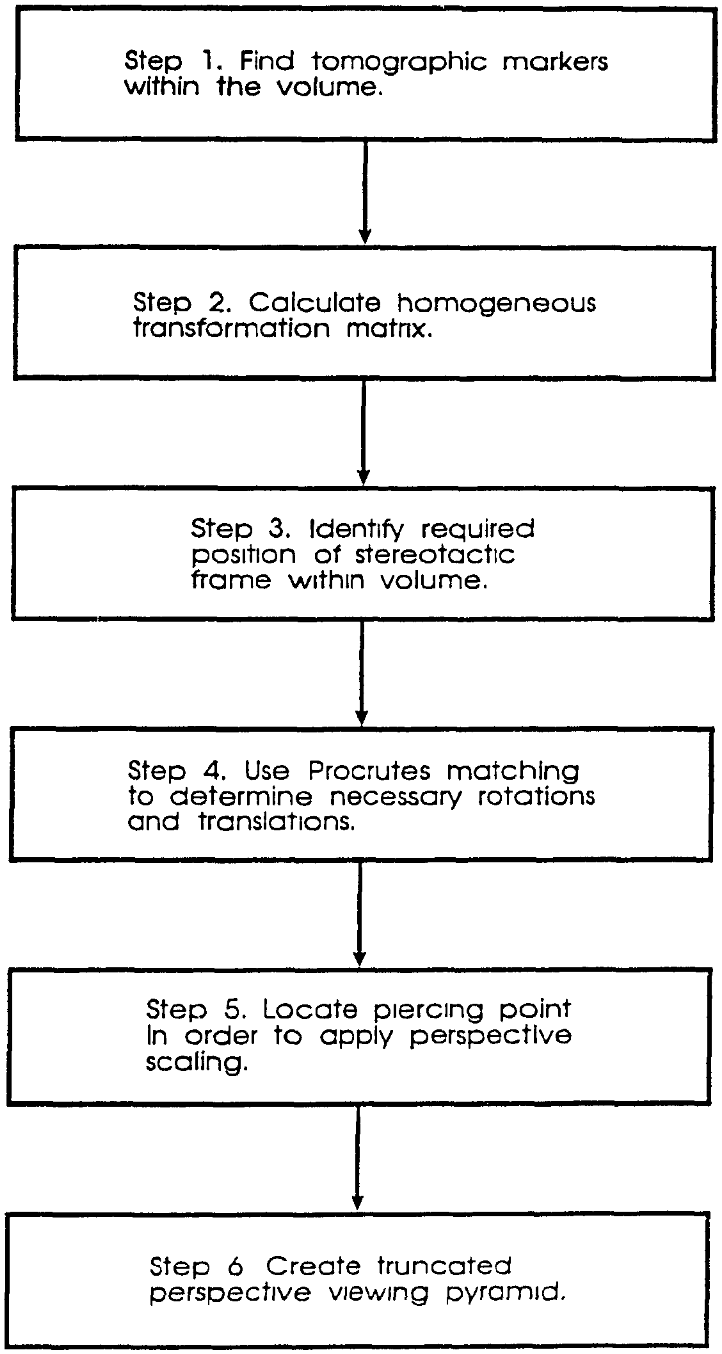

Figure 4.9: Manipulation of volumetric data to yield perspective projection.

set has the correct orientation and position relative to the image plane (the projection plane).

5. The piercing point P must be located in voxel coordinates before applying the perspective scaling which is given by the intersection of the viewing ray with the image plane. The column of voxels perpendicular to the projection plane and containing P is equivalent to the principal viewing ray and all perspective scaling is centered on this column.

6. In order to obtain a perspective rendering when using orthogonal projection, it is necessary to resize the volume slices that are parallel to the projection plane. The result of this operation is the perspective viewing pyramid[1]. The following algorithm is applied to each slice:

I take point J on viewing ray

II transform J_{voxel} to J_{frame}

III let $J'_{screen} = [M]\, J_{frame}$

IV choose point I at the edge of slice on the same horizontal scan line as J

V transform I_{voxel} to I_{frame}

VI let $I_{screen} = [M]\, I_{frame}$

VII choose point K at the edge of slice on the same vertical scan line as J

VIII transform K_{voxel} to K_{frame}

IX let $K_{screen} = [M]\, K_{frame}$

X let $C_x = |Ix_{screen} - Jx_{screen}| / |Ix_{frame} - Jx_{frame}|$

XI let $C_y = |Ky_{screen} - Jy_{screen}| / |Ky_{frame} - Jy_{frame}|$

[1]One should note that the sides of the pyramid are not flat, but are curved since perspective magnification is not linear

XII C_x and C_y are the scaling factors to be applied to the slice.

After the perspective viewing pyramid is formed, the direct volume rendering algorithm is applied to create the required perspective rendering.

4.4 Results

Figure 4.10 shows a single frame from a sequence of 32 images showing a the rotation of a semi-transparent rendering of a 64 slice MR data set with a posterior section removed. The transparency is more apparent in the animated sequence. The tissue is segmented using an intensity-based heated-object look-up table (LUT), which maps (starting from the lowest intensities) to black, then red, through yellow and then white. This scheme enhances contrast without creating artificial contours seen in rainbow colour scales. Surface enhancement is achieved by matting the classified volume against the gradient magnitude before the data set is projected onto the image plane using the standard PIXAR technique.

Figure 4.11 shows the same data set rendered with the modified algorithm using a single viewer centered light source. This image demonstrates that it is easier to distinguish differences in depth with the more opaque rendering.

Figure 4.12 shows the rendering of two 64 slice MR angiograms of the neck. The vessels are segmented from each volume based on their pixel intensity by mapping through a linear continuous classification LUT. In this example, the arterial volume is coloured red and the venous blue, with the vessels in both volumes being assigned high opacities relative to their respective pixel intensities. A composite volume is formed by direct addition of the two classified data sets without surface enhancement. The projections were then calculated by simple depth-cued opacity-integrals. As can be seen, the technique yields projections in which objects appear semi-transparent, giving some depth cues through partial occlusion. The depth shading, achieved using a ramped matte volume, reduces ambiguity in the visualization often apparent as a reversal in the

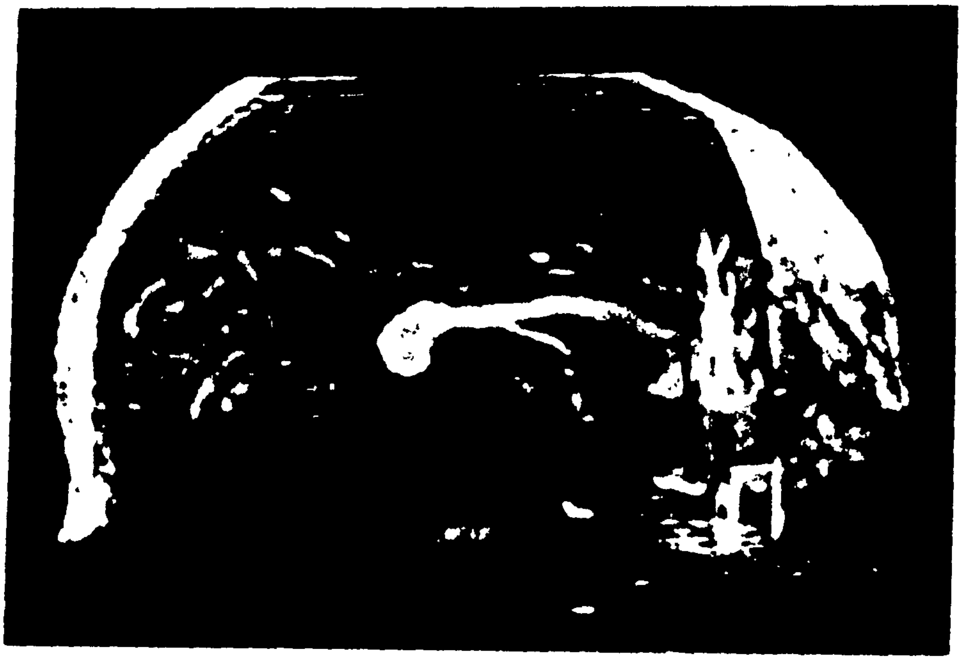

Figure 4.10: Semi-transparent rendering of 64 slice MR data set using the PIXAR technique.

Figure 4.11: Rendering of 64 slice MR data set using the modified algorithm.

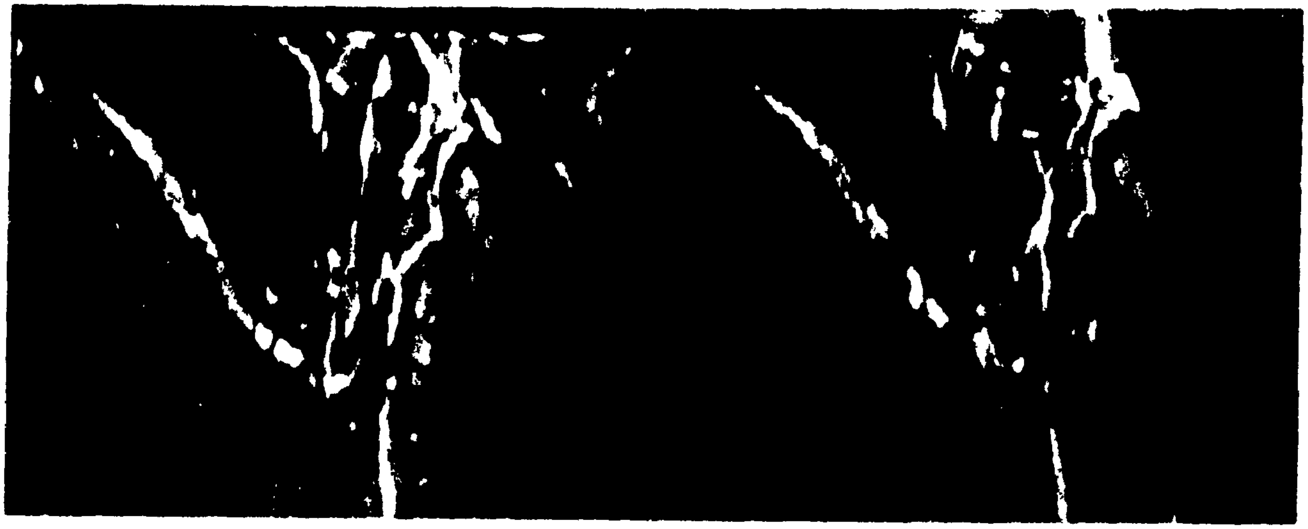

Figure 4.12: Semi-transparent rendering of composite MR angiogram of arteries (red) and veins (blue) obtained as two 60 slice acquisitions and rendered using the PIXAR depth-cued opacity-integral technique.

direction of rotation or as an inside-out flip in depth.

Figure 4.13 shows an example of combined vasculature and cortical data acquired as a single MRA study. To expose the surface of the brain, the cortical tissue was outlined using a semi-automated intensity contouring procedure. The contours were then filled and convolved with a gaussian kernel with a full-width at half-maximum (FWHM) of 5 pixels to produce a blurred a matte volume used to mask the cortex from the surrounding tissue prior to projection. This blurring operation is necessary to ensure a gradual change at the mask edge in order to avoid artifacts in the final rendering due to artificially high gradient values at a binary mask boundary. The masked cortex was then mapped though a heated-object colour LUT. Vessels were extracted as described above, based on their high pixel intensity and mapped to white. The vessels were then composited *OVER* the cortical tissue. The opacity of both the cortex and the vessels were

Figure 4.13: Volume rendering of cortical surface and vasculature. The data was acquired as a single 128 slice acquisition and processed using the modified algorithm.

Figure 4.14: Merged rendering of MR tomographic data set and corresponding DSA image. This stereo-pair may be viewed using the crossed-eye technique.

set to high values based on their original intensities and the projections were formed with the modified algorithm, using the gradient calculated from the original volume as estimates of voxel surface normals.

Figure 4.14 is an example of the merging technique applied to DSA images and MR tomographic data and Figure 4.15 is the result when applied to a CT volume. In both cases, the fiducial markers are identified in an unsubtracted angiogram and used to calculate the DSA image geometry and the necessary volume manipulation parameters. The tomographic volumes are then loaded into the PIXAR, interpolated to cubic voxels, oriented to match the required view and differentially scaled to match the perspective magnification (as described in detail in §4.3.3). The MR volume is rendered using the modified algorithm (shown in Fig. 4.14), while the CT volume is rendered using the original PIXAR method (shown in Fig. 4.15).

Figure 4.15: Merged rendering of CT data set and corresponding DSA image. This stereo pair may be viewed using the crossed-eye technique.

4.5 Problems with Rendering and Merging

4.5.1 Robustness to parameter selection

Images created by the rendering algorithm are sensitive to the parameters chosen for opacity mapping in the segmentation algorithm and to the parameters chosen for lighting direction. A small change in the opacity mapping in the segmentation has a large affect on surface visualization. A steep opacity ramp at the edge of a region results in a very hard, noisy surface. However, a ramp that climbs too slowly causes definition in the cortical folds to be lost, resulting in an image with low contrast surface and details.

The lighting model used by the ChapVolume software library is a black box. It is limited in that it is not possible to control the amount of diffuse and specular light used to illuminate a model. However, if a single light source is used, centering the light source at the viewing focal-point (viewer-centered lighting) yields acceptable images. Multiple light sources can be simulated by adding the enhancement due to each source to the 2D slice before projection.

4.5.2 Dependency on segmentation algorithm

The first step in the rendering algorithm consists of mapping the acquired data into a matte volume used for segmentation. Each voxel in the input volume is essentially classified using a Bayesian estimate of the percentage of each tissue of interest within that voxel, since a binary classification will introduce artefacts in the final projection rendering. The use of a probabilistic classifier reduces the dependency on the actual results of the segmentation. However, a misclassified volume is not interpretable. Future projects will need to include work on better algorithms for automatic segmentation.

4.5.3 Robustness to slice thickness

Direct 3D volumetric rendering works best when the data set has isotropic resolution, sampled at (or above) the Nyquist frequency. When this is not the case, aliasing artifacts are apparent in the rendered image such as stair-casing, banding or striping. When the slice thickness is large compared to the inplane resolution, surface rendering techniques (marching cubes, dividing cubes, surface contour tesselation) produce more attractive images. However, I am not certain that this is better, since it gives the viewer a false sense of data quality. Work has been done in the context of intelligent interpolation by Liang et al. in [55], in order to overcome this problem for volumetric rendering. Their technique use a contour following algorithm after structures have been identified in two consecutive slices. To create an slice between the two originals, the contours and the grey level regions contained within them, are interpolated. Unfortunately, this technique depends on a good segmentation algorithm. A more direct solution would involve modification to scanning protocols to obtain volumes with isotropic resolution.

4.5.4 Occlusion artifacts of merged DSA on tomographic volume

The rendered images contain a number of depth cues to give the impression of a 3D scene. Among these include depth shading, perspective, occlusion and lighting highlights. When a pair of images is rendered to match the stereo DSA images, stereo cues also aid in the perception of depth. However, the such cues will oppose those of occlusion when a vessel apparently enters an opaque tissue and is not hidden by it. By making transparent renderings of the of the anatomical surfaces, the depth cues due to occlusion are reduced, and consequently minimize possible conflicts.

4.5.5 Artifacts of pseudo stereoscopy

The stereoscopic projection images are rendered from a particular viewpoint. If the images are not viewed from the same viewpoint, such as when the neurosurgeons flip left and right images when viewing with crossed-eye stereo, the rendered depth cues conflict with the stereo cues. Simple additive projections could be used for both forward and backward viewing, however a simple radiograph would give the same information with a much higher quality and volume rendering would not be necessary.

4.5.6 Computational cost

Direct volume visualization is a computationally expensive technique to make 3D projection images. Interactive manipulation of the volume would be necessary (but currently impossible) to examine all of the data. Since the renderings are merged with given stereo views, the possible viewing geometry is limited to specific AP and lateral images. It is possible to pre-calculate a sequence of images (each with a different amount of tissue) for interactive panning through the volume. With a small amount of pre-planning of the visualization, thirty to sixty images could be rendered in the hour before the surgeon sits down for the analysis and planning.

4.6 Summary

The rendering techniques presented in this chapter produce projections of tomographic data sets that allow the user to visualize the anatomy in a natural 3D format. The renderings do not use any specific surface detection schemes, rather natural surfaces are enhanced during the rendering operation with the task of interpretation left to the observer.

The merging of the DSA projection with the tomographic volume depends on the use of the fiducial markers and the stereotactic frame. It may be possible to use anatomical knowledge, locating a number of anatomical landmarks in both the volume and the projection, along with the known scanning parameters of both modalities, to create a matched image.

A large amount of the information in the tomographic volume is correlated. The shaded surface renderings enhance the surfaces existing the in data, and in a sense may show all of the relevant information from the study. If only the matched stereo views of the 3D tomographic reconstruction and DSA projection described above need to be saved, the rendering technique may be considered as an effective, destructive, high-rate data compression scheme. While the tomographic data sets require large memory areas for storage, the resulting stereo pair occupies a fraction of the space.

The images created contain information of both vasculature and anatomy Since the renderings are created using the same geometry as the DSA projection, the existing stereoscopic stereotactic analysis software may be used to locate points, measure distances and plan probe trajectories in the anatomical volume. This technique for combining data from different modalities is new and constitutes the first step towards the realization of a true 3D stereotactic surgery planning workstation It is envisaged that in its ultimate realization, the surgeon will interact with the image using a hand-held 3D localizing probe.

Chapter 5

Projection Accuracy

Rendered images of tomographic data are merged with DSA images for quantitative use in the stereotactic planning system. Rendered image quality depends on a number of factors including noise, contrast, spatial distortion and sharpness. (A good description of each cause is given by Hendee in [34].) This chapter addresses the question of image sharpness; a measure of the blurring of an object's true projection represented in a rendered image. The blurring is due to original data sampling limits, patient motion artifacts, image rendering geometry, and finally the display characteristics.

Sampling and patient motion artifacts are not treated here, as the use of the stereotactic frame (in CT and MR) or foam-lined head-holder (for PET) fixed to the scanner couch minimizes most motion artifacts. Assumptions made in chapter 4 regarding sampling apply here. That is, the volumetric data set, $\mathcal{V}$, represents an object, $\mathcal{O}$, sampled in 3D space at or greater than the Nyquist frequency. If this is not possible, the assumption is made that the value associated with a voxel represents the average measured value within the volume element, which represents a low-pass filtering of the object before sampling.

The image quality can be characterized quantitatively in two ways. The ideal projection of a point, $\mathbf{a}$, from the object, $\mathcal{O}$, is described by $\mathcal{P}(\mathbf{a})$. If $P()$ is the rendering projection operator, then the point spread function (PSF) of the volume rendering process is defined by the projection of a single voxel, $P(\mathbf{a})$, onto the image plane. Similarly, the modulation transfer function (MTF) describes the blurring by specifying how well each spatial frequency in the true projection, $\mathcal{P}(a)$, is preserved in the volume rendering projection, $P(V)$. The MTF is the normalized modulus of the Fourier transform of the PSF.

This chapter begins by deriving the PSF for straight orthographic projections, followed by arbitrary angle orthogonal views and arbitrary perspective views. The results are confirmed experimentally and the MTF is measured by computer rendered images of random points distributed in the data volume.

5.1 Orthogonal projection

When creating a direct orthogonal projection the PSF is due to the projection operator only. Suppose $\rho(x,y,z)$ is the density function of the object, $\mathcal{O}$, convolved with the imaging system PSF. Let the $x-$, $y-$ and $z-$axis define the world (or stereotactic frame) coordinate system. The sampled density function, $\rho_s(x,y,z)$, is defined by measurements performed at a finite sampling rate described by:

$$\rho_s(x,y,z) = comb(\frac{x}{x_s}, \frac{y}{y_s}, \frac{z}{y_s})\rho(x,y,z) \tag{5.1}$$

where x_s, y_s and z_s are the sample intervals in the $x-$, $y-$ and $z-$directions, respectively of the Dirac delta functions of the 3D *comb* function [7], with the resulting samples being stored in the volumetric data set, $\mathcal{V}$. The $X-$, $Y-$ and $Z-$axes define the volume coordinate system, with transformation matrix T_{wv} defining the world-to-voxel coordinate system mapping. The sampling volume of $\mathcal{O}$ is limited in space by x_{min}, x_{max}; y_{min}, y_{max}; z_{min}, z_{max} in the volume coordinate system.

Images are formed by projecting the volumetric data set onto the image plane, whose coordinate system is defined by the $u-$ and $v-$axis at $w = 0$. The $w-$axis extends from of the image plane towards the viewer. The transformation matrix T_{vs} defines the voxel to screen mapping.

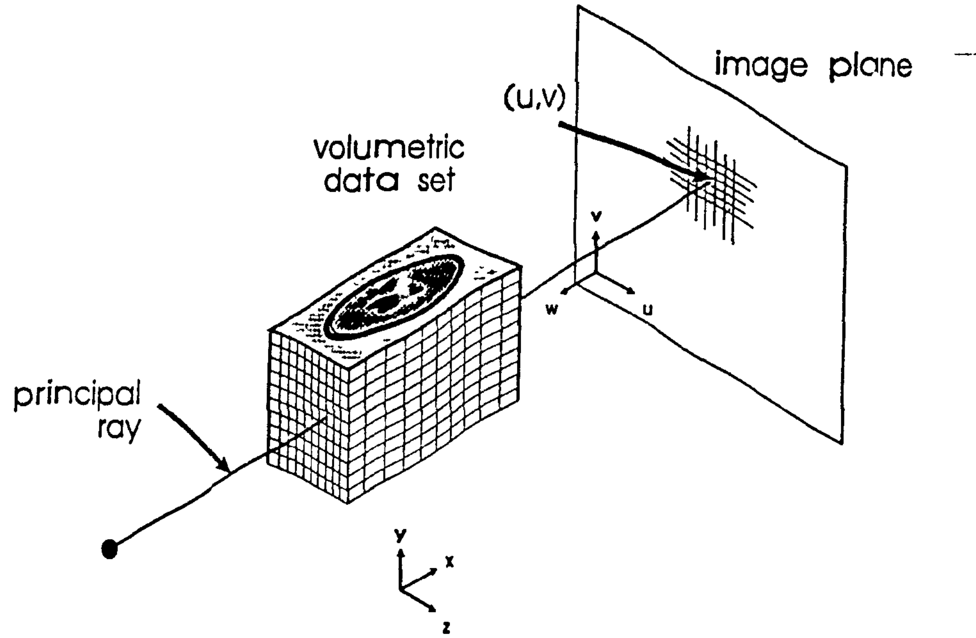

Figure 5.1: Orthographic projection of the volume. The projection rays are parallel to a major axis of the volume. The rays pass through the volume, resampling it, and project onto the image plane.

When creating a projection image, it is necessary to sample the measured data volume arbitrarily along a given ray (see Fig. 5.1). An estimate of the continuous function, $\rho(x, y, z)$, within the volume must be reconstructed from $\rho_s()$ prior to resampling. If the object, $\mathcal{O}$, is spatially band-limited (ρ has negligible frequency components outside the range $|f_x| \leq B_x$, $|f_y| \leq B_y$, $|f_z| \leq B_z$, where f_x, f_y, f_z are the signal frequencies and B_x, B_y, B_z are the frequency limits in the $x-$, $y-$ and $z-$axis respectively) and sampled at a rate greater than the Nyquist frequency, its density function, ρ, can be exactly reconstructed from the sampled data points by *sinc* interpolation [68]:

$$
\begin{aligned}
\rho &= \rho_s * h_V, \qquad (5.2)\\
&= \rho_s * [sinc(2B_x x) sinc(2B_y y) sinc(2B_z z)],\\
&= \int_{x_{min}}^{x_{max}} \int_{y_{min}}^{y_{max}} \int_{z_{min}}^{z_{max}} \rho_s(x', y', z') \left[\frac{sinc(x - 2B_x x')}{2B_x} \frac{sinc(y - 2B_y y')}{2B_y} \frac{sinc(z - 2B_z z')}{2B_z} \right] dx\, dy\, dz,\\
&= \frac{1}{8 B_x B_y B_z} \sum_{p=z_{min}}^{z_{max}} \sum_{n=y_{min}}^{y_{max}} \sum_{m=x_{min}}^{x_{max}} \rho_s(m, n, p) sinc(x - 2B_x m) sinc(y - 2B_y n) sinc(z - 2B_z p),
\end{aligned}
$$

where '$*$' signifies the 3D convolution operator, h_V is the separable 3D reconstruction kernel and x', y' and z' are the variables used for integration in the continuous domain and m, n and p are the corresponding variables used in the discrete domain. In a simplified form[1]

$$\rho(\mathbf{x}) = \sum_{\mathbf{i} \in Vol} \rho_s(\mathbf{i}) h_V(\mathbf{x} - \mathbf{i}), \tag{5.3}$$

where $\mathbf{i} = (i_x, i_y, i_z)$ is a sample point and the summation over Vol covers all the samples within the extent of the volumetric reconstruction kernel h_V centered on $\mathbf{x} = (x, y, z)$. *Sinc* interpolation is used in this derivation since it appears[2] that the PIXAR implementation of the affine transformations used to geometrically manipulate the volume uses a truncated sinc kernel and the use of a sinc reconstruction kernel greatly simplifies the derivation.

The value of points $\mathbf{u} = (u, v)$ in the image plane is calculated by integration along a line through the volume (see Fig. 5.1). Since orthographic projections are used, the rays are parallel to the $Z-$axis and the $(x; u)$ and $(y; v)$ coordinate systems are aligned. Due to the imaging geometry, the rays are also parallel to the $w-$axis of the image plane coordinate system. For an additive projection image, the point values in the image plane are given by:

$$\begin{aligned} I(\mathbf{u}) &= \int_{z_1}^{z_2} \rho(\mathbf{w}) dz, & (5.4) \\ &= \int_{z_1}^{z_2} \sum_{\mathbf{i} \in Vol} \rho_s(\mathbf{i}) h_V(\mathbf{w} - \mathbf{i}) dz, & (5.5) \end{aligned}$$

where z_1 and z_2 are the minimum and maximum volume limits in the $Z-$direction and $\mathbf{w} = (u, v, z)$.

Instead of reconstructing an arbitrary point by convolving the reconstruction kernel with all the sample points and projecting the result to the image plane, consider projecting all the samples to the image plane and convolving them with the projection

[1] Using notation from Westover in [101]

[2] The affine transformations implemented by PIXAR are proprietary, and the documentation does not explicitly state the type of kernel used. However, experiments discussed at the end of this chapter indicate the use of a *sinc* kernel, truncated after the first negative lobe

of the reconstruction kernel. Thus, the reconstruction of the projection of an arbitrary point may be performed completely in the image plane.

If the summation over $\mathbf{i}$ is moved out of the integral in eq. 5.5, then the intensity of points in the image plane is due to the sum of the projection of each sample point's contribution to the reconstruction of a given point $w = (x, y, z)$.

$$I(\mathbf{u}) = \sum_{\mathbf{i} \in Vol} \int_{z_1}^{z_2} \rho_s(\mathbf{i}) h_V(\mathbf{w} - \mathbf{i}) dz. \tag{5.6}$$

Consider how a single sample point, $\mathbf{i}$, contributes to the reconstruction of an arbitrary point $\mathbf{x} = (x, y, z)$. The point $\mathbf{i}$ must be included within the extent of the volume reconstruction kernel centered on $\mathbf{x}$. This scalar value $\tau_{\mathbf{i}}$ is given by:

$$\tau_{\mathbf{i}}(\mathbf{x}) = \rho_s(\mathbf{i}) h_V(\mathbf{x} - \mathbf{i}). \tag{5.7}$$

The projection intensity due to sample $\mathbf{i}$ is given by:

$$\begin{aligned} I_{\mathbf{i}}(\mathbf{u}) &= \int_{z_1}^{z_2} \tau_{\mathbf{i}}(\mathbf{x}) dz, & (5.8) \\ &= \int_{z_1}^{z_2} \rho_s(\mathbf{i}) h_V(\mathbf{x} - \mathbf{i}) dz. & (5.9) \end{aligned}$$

Since the sample point $\rho_s(\mathbf{i})$ is a constant independent of z, it can be moved out of the integral:

$$I_{\mathbf{i}}(\mathbf{u}) = \rho_s(\mathbf{i}) \int_{z_1}^{z_2} h_V(\mathbf{x} - \mathbf{i}) dz. \tag{5.10}$$

Substituting eq. 5.10 in eq. 5.5 gives:

$$I(\mathbf{u}) = \sum_{\mathbf{i} \in Vol} \rho_s(\mathbf{i}) \int_{z_1}^{z_2} h_V(\mathbf{x} - \mathbf{i}) dz, \tag{5.11}$$

i.e., the intensity field in the image plane due to the projection of an arbitrary point, $(\mathbf{x})$, is equal to the sum over the convolution of the reconstruction kernel's projection with each sample point within the extent of the kernel. Replacing $h_V()$ with the separable 3D *sinc* interpolation function gives:

$$I(\mathbf{u}) = \sum_{\mathbf{i} \in Vol} \rho_s(\mathbf{i}) sinc(x - i_x) sinc(y - i_y) \int_{z_1}^{z_2} sinc(z - i_z) dz. \tag{5.12}$$

If z_1 and z_2 represent the minimum and maximum limits in the z-direction in the sampled volume then $\int_{z_1}^{z_2} sinc(z - i_z)dz \approx 1$ and the image plane intensity field is due to 2D *sinc* interpolation of the projected sample points. The intensity of the point spread function of an arbitrary point $\mathbf{x}$ that projects to the image point $\mathbf{u}$ (i.e. $(u, v) = T_{vs} \cdot (x, y, z)$, where T_{vs} is the voxel to screen transformation) is defined by:

$$I(\mathbf{u}) = \sum_{\mathbf{i} \in Vol} \rho_s(\mathbf{i}) sinc(x - i_x) sinc(y - i_y) \tag{5.13}$$

If the object has been sampled at a rate exactly equal to the Nyquist frequency, then the zero-crossings of the 3D *sinc* function will fall at sample points (i.e. the value of the reconstruction kernel for a given sample point will be equal to zero at the position of the neighboring points), so that the reconstruction of a point coinciding with a sample point is the sample itself. If the image is projected onto a matrix with the same sampling rate and phase as the planes parallel to it in the volume (i.e. the pixels line up exactly with the voxels), then the projection of that data point is simply a voxel-to-pixel copy and add. This is exactly the method used for projection rendering by the PIXAR's ChapVolume software library.

For an orthographic projection the reconstruction kernel does not change for different voxels, hence the point spread function is spatially invariant [101]. However in a perspective view, the sampling rate changes with respect to distance from the screen Therefore, the PSF function must be modified to take into account the divergent rays of a perspective projection, or in the case of the truncated viewing pyramid, the refinement must take into account the changing sampling rate.

5.2 Arbitrary orthogonal projection

Since orthographic projections are used by the PIXAR to render images, the sampled volume must be transformed to match the required view. An arbitrary orthogonal view of $\mathcal{O}$ is formed by applying rotation ($\mathbf{R}$) and translation ($\mathbf{T}$) transformations to $\mathcal{V}$ before projection, so that the front face of the volume is parallel to the image plane. If $\mathcal{O}$ is sampled isotropically and the measurement of the density function satisfies the Nyquist criteria then the transformations are equivalent to resampling V by a 3D array that is positioned and oriented to give the required view by $\mathcal{V}'$.

If the density function is band-limited (with bandwidth B) then the convolution given by eq 5.2 becomes:

$$\begin{aligned} \rho &= \rho_s * h_V, & (5.14) \\ &= \frac{1}{8B^3} \sum_{p=z_{min}}^{z_{max}} \sum_{n=y_{min}}^{y_{max}} \sum_{m=x_{min}}^{x_{max}} \rho_s(m,n,p) sinc(x - 2Bm) sinc(y - 2Bn) sinc(z - 2Bp). \end{aligned}$$

where x_{min}, x_{max}, y_{min}, y_{max}, z_{min}, z_{max}, m, n, p, x, y and z are as defined for eq. 5.2. The density function is exactly reconstructed from its sampled points prior to resampling with the new 3D array to form $\rho'_s()$. This array must also fulfill the criteria for Nyquist sampling frequency (i.e. have the same spacing as in the first array) so that arbitrary points in the transformed volume may be reconstructed.

Equation 5.3 gives the exact value for the density function ρ at any point $\mathbf{x}$ in the original data volume. The density value of all sample points $\mathbf{i}'$ in the resampled volume are calculated by applying the inverse transformation, $\mathbf{T}^{-1}\mathbf{R}^{-1}$, to $\mathbf{x}'$ to find the position of the point $\mathbf{x}$ in the original volume. Equation 5.3 is then used to give the value for $\rho(\mathbf{x})$ which is equal to $\rho'_s(\mathbf{x}')$. The density of a point in the transformed volume is given by:

$$\begin{aligned} \rho' &= \rho'_s * h_V, & (5.15) \\ &= \frac{1}{8B^3} \sum_{p=z'_{min}}^{z'_{max}} \sum_{n=y'_{min}}^{y'_{max}} \sum_{m=x'_{min}}^{x'_{max}} \rho'_s(m',n',p') sinc(x' - 2Bm') sinc(y' - 2Bn') sinc(z' - 2Bp'). \end{aligned}$$

where x'_{min}, x'_{max}, y'_{min}, y'_{max}, z'_{min}, z'_{max}, m', n', p', x', y' and z' correspond to their unprimed counterparts as defined for eq. 5.2.

The PSF of the projection of an arbitrary point is given by applying eq. 5.13 to the point $\mathbf{x}'$ using $\rho'_s()$ instead of $\rho_s()$ and integrating along the z'-axis in the reoriented volume. If the volume $\mathcal{V}'$ is isotropically sampled at greater than Nyquist, the projection of the 3D kernel will have the same shape, regardless of the projection direction.

5.3 Perspective projection

A perspective projection is created by scaling each slice in the volume parallel to the image plane proportional to its distance to the plane. This creates a truncated perspective viewing pyramid, $\mathcal{V}''$, with its base parallel to the image plane. While $\mathcal{V}'$ continues to have isotropic resolution, $\mathcal{V}''$ is no longer isotropically sampled, since the sampling rate changes when the slices are scaled. If $\mathcal{V}'$ is sampled exactly at the Nyquist limit, the size of the slices can only be increased since scaling them down violates the sampling criteria and arbitrary points can no longer be recovered from the sampled points.

To calculate the reconstruction kernel of an arbitrary point $\mathbf{x}'' = (x'', y'', z'')$ in the viewing pyramid, its corresponding source point $\mathbf{x}' = (x', y', z')$ in the oriented volume must be found. This is done using the scaling factors that were used to create the truncated perspective viewing pyramid.

The point $\mathbf{x}'$ has a separable 3D *sinc* reconstruction kernel associated with it as shown in the last section. (The 3D shape at FWHM of the its central peak is a right hexahedron, i.e. *box-shaped.*) The kernel associated with $\mathbf{x}''$ is estimated by geometrically scaling the one of $\mathbf{x}'$ by the same parameters used to scale the slice containing $\mathbf{x}'$. It will be scaled in the $x'-$ and $y'-$directions.

The reconstruction kernel of $\mathbf{x}''$ can be calculated more exactly by differentially scaling $\mathbf{x}''$s kernel from front to back. The resulting shape of its peak at FWHM resembles a trapezohedron (truncated pyramid).

The resulting point spread function is calculated by projecting the reconstruction kernel onto the image plane, given by:

$$I(u,v) = \sum_{\mathbf{i}' \in Vol} \rho'_s(\mathbf{i}') sinc(\frac{x' - i'_x}{S_{x'}(i'_z)}) sinc(\frac{y' - i'_y}{S_{y'}(i'_z)}) \tag{5.16}$$

where $S_{x'}(i'_z)$ and $S_{y'}(i'_z)$ are the perspective scaling parameters in the x'- and y'-directions.

5.4 Practical limits

The preceding discussion is based on two assumptions: all sampling is performed isotropically and at a rate at least twice that of the highest frequency component in the original density function; and the reconstruction filter is designed to pass all frequency components of the reconstructed point without distortion. Pratt [81] defines two contributions to the reconstruction error; the difference between the interpolation function used and ideal reconstruction kernel, and the finite bounds of the reconstruction kernel that cause truncation artifacts. Since the actual implemented reconstruction kernel is limited in size, arbitrary points are reconstructed to finite precision.

A number of interpolation kernels can be chosen. A square pulse function is the easiest to implement and results in a nearest neighbour or zeroth-order interpolation. A triangular pulse is used for linear interpolation (and is the result of the convolution of two square waves). The reconstruction kernel used in the affine transformations of the pixel windows in the ChapVolume software library [2] is the truncated sinc function.

Each resampling operation can be seen as flattening the response from a point source by convolution of each point in the volume by the PSF. In order to create the perspective viewing pyramid there are seven resampling operations. The final image is convolved with six point spread functions before the perspective scaling is applied. However, not all six are applied in the same direction, for example a rotation about one axis blurs the data in each plane perpendicular to the rotation axis. The blurring is considered to have two components, one along each in-plane axis due to the two resampling steps applied during the rotation algorithm. Hence, the three rotations contribute two convolutions to each axis. The translations blur the data only in the direction of the translation, adding one convolution to each axis. Therefore, the extent of resulting PSF is expected to be three times broader than that due to a single convolution (before perspective scaling is applied). However, its FWHM increases by only 80approximated by the convolution of three gaussian kernels. Thus, the expected resolution is on the order of 1.8 pixels.

5.5 Experiments

In the previous section, a PSF was defined theoretically in terms of the affine transformation interpolation filter used to geometrically manipulate the volume before projection. Here, experiments are performed using two methods to confirm the estimates made above. First the PSF determined by rendering a single voxel is plotted and, second, the MTF of the projection of a random field is measured.

In order to separate scanning and imaging errors from those involved with the rendering, the data sets for the experiments are created mathematically. Single, randomly placed voxels are used to measure the PSF and a uniform random field placed within the volume is used to determine the MTF. The volumes are rendered using the direct volume rendering technique described in Chapter 4.

The MTF shows how the unit spectrum is modulated by the system. The MTF curve describes what fraction of the input is transferred though the system at each frequency, detailing the system's ability to transfer the lowest possible sinusoidal contrast differences. For a given object contrast, O_c, and image contrast, I_c,

$$O_c = \frac{S_{max} - S_{min}}{S_{max} + S_{min}}, I_c = \frac{I_{max} - I_{min}}{I_{max} + I_{min}}, \tag{5.17}$$

where S_{max} and S_{min} are the maximum and minimum object intensities and I_{max} and I_{min} are the maximum and minimum image intensities, the MTF is the ratio of image contrast to object contrast for a given sinusoidal frequency of contrast change, and when is evaluated at each frequency, we have

$$MTF(f) = I_c(f)/O_c(f). \tag{5.18}$$

Direct measurement of the MTF is tedious, however it may be determined by calculating the modulus of the Fourier transform of the PSF,

$$MTF(f) = \frac{|\int_{-\infty}^{\infty} p(x)e^{-2\pi i f x}dx|}{\int_{-\infty}^{\infty} p(x)dx}, \tag{5.19}$$

where $p(x)$ is the amplitude of the PSF, measured at a distance x from its center or by calculating the ratio of the power spectrum of the system output (the projection image) over the input (the random field).

In the spatial domain, the PSF is the 2D unit response function, which describes how an input point is transformed by the system. In general, it may or may not be spatially invariant. A single pixel in the spatial domain is represented by all frequencies in the spectrum. Since the projection operator is not ideal, the projection of a voxel to the image plane is not represented by a single pixel, but rather by a small group of pixels and the MTF shows characteristics of a low-pass filter. Since the image is projected onto a square matrix, the PSF (and consequently the MTF) is asymmetric. In the invariant case (i.e. orthogonal projection), the 2D MTF can be calculated directly from the 2D PSF, just as in eq. 5.19 giving:

$$MTF(f_x, f_y) = \frac{\left| \int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} p(x,y) e^{-2\pi i f_x f_y x y} dx dy \right|}{\int_{x=-\infty}^{\infty} \int_{y=-\infty}^{\infty} p(x,y) dx dy} . \tag{5.20}$$

In the following tests of the first experiment, the PSF is determined by averaging the projection of 10 different randomly placed voxels within the volumetric data set. Each voxel is set to full intensity (=2047). Geometric transformations, similar to those required to prepare a volume for a given view, are applied. (Perspective magnification is not applied, as it is described by a simple scaling factor.) The volume is then projected using the additive projection technique described in chapter 4. A program is used to automatically locate each voxel's projection, and average them after applying a cross-correlation to match their centres.

The first example, figure 5.2, shows a 2D elevation plot, the vertical profile and the horizontal profile of the average PSF of an orthographically projected voxel subjected to translations in the $x-$ and $z-$directions in order to position the volume in front of the image plane. The spread along the $x-$axis is easily seen projected onto the $u-$axis of the image plane. The spread in the $z-$direction is integrated along the $w-$axis, essentially nullifying the blurring effect along this axis.

The next three tests demonstrate how the PSF changes with rotations about each axis after the application of a translation operator. Figure 5.3 shows the average PSF

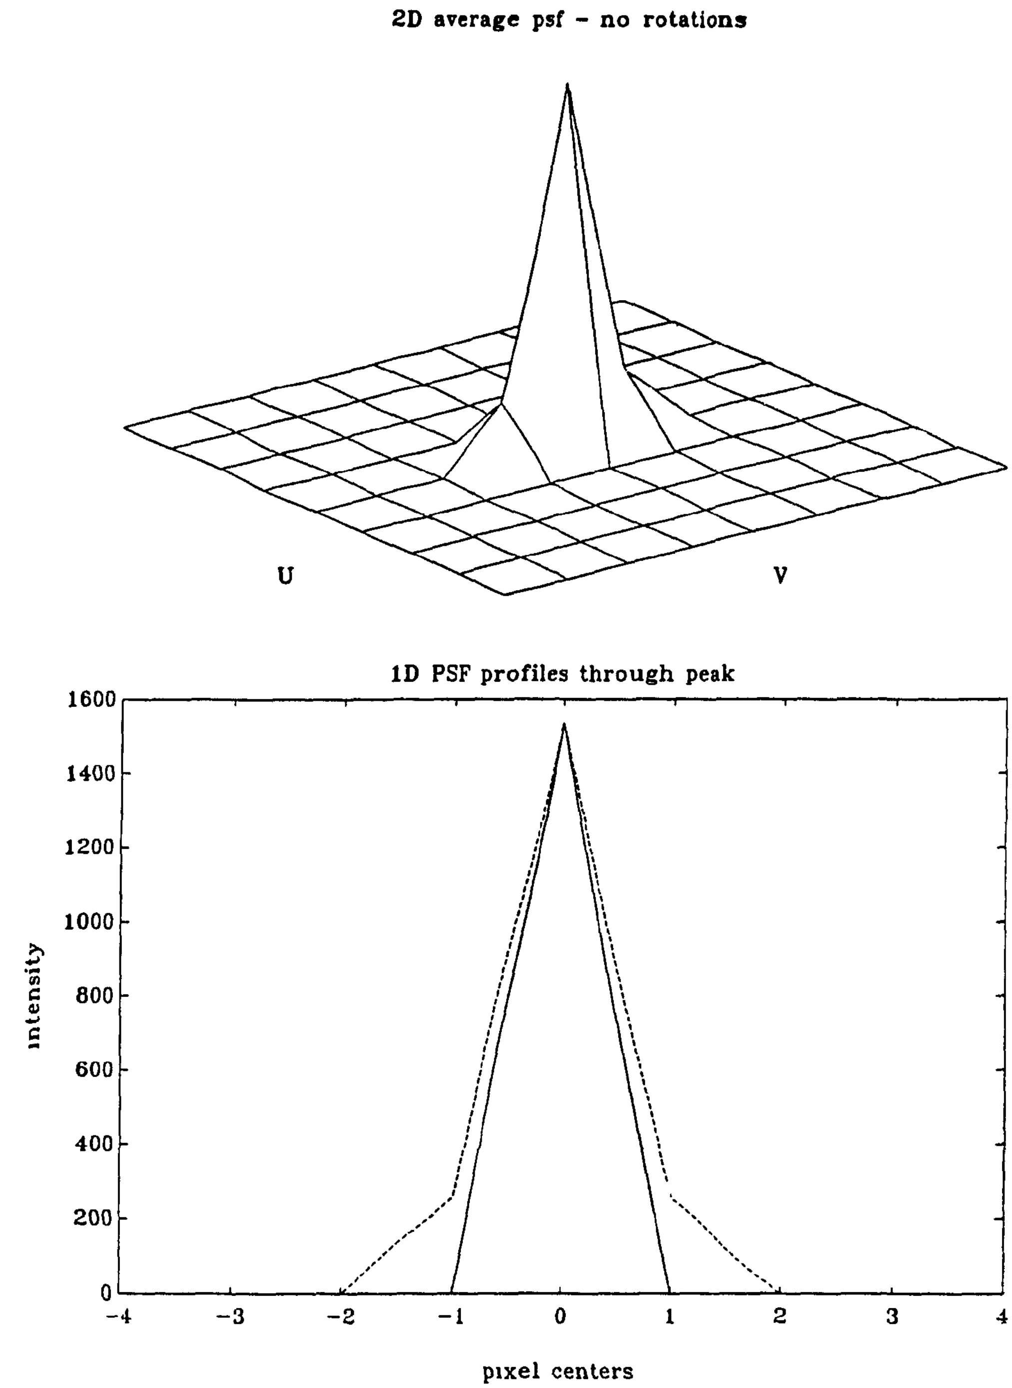

Figure 5.2: PSF of voxel after translations only. a) shows 2D elevation plot of PSF, b) shows horizontal and vertical profiles through the PSF peak.

of an orthographically projected voxel rotated about the $x-$axis after translation. The rotation blurs the data in the $yz-$plane. The projection integrates the blur once again along the $w-$axis, resulting in a PSF that is slightly wider in the $v-$direction than that due to translation only.

Figure 5.4 shows the average PSF of orthographically projected voxels rotated about the $y-$axis after translations. The blur in the $xz-$plane is projected onto the $u-$axis of the image plane.

Figure 5.5 shows the average PSF resulting from a rotation about the $z-$axis after translations for orthographic projections. The blur in the $xy-$plane is projected directly onto the $uv-$plane and it easily seen that the PSF is symmetrically wider than the PSF corresponding to translation only.

The fourth test, shown in Fig. 5.6, examines how the PSF changes with rotation angle. A voxel in the volume is rotated about the $y-$axis and projected to a single scanline. The 1D profile of the projected point is measured and a gaussian is fit to the sample points in order to approximate the width of the PSF. The FWHM and full-width-tenth-maximum (FWTM) of the gaussian curve are measured and plotted for rotation angles from zero to ninety degrees at 1 degree increments. The measured FWHM increases from 1.0 to 1.5 pixels as the angle changes from zero to forty-five degrees. The width of the FWHM then diminishes back to one pixel as the angle increases to ninety degrees. This is due to the PIXAR implementation of the Catmull-Rom [10] rotation algorithm. Since the PSF increases with the the angle of rotation, transposition and reflection of the pixel window using direct pixel-to-pixel copies are used to limit the necessary rotation angle to the range -45^o to 45^o, thus ensuring a minimum spread.

Finally, figure 5.7 shows the average PSF from an orthographically projected voxel after translations and rotations about all three axes. The FWHM of the PSF is approximately 2 pixels wide in both the horizontal and vertical directions. The corresponding 2D MTF is shown in Fig. 5.8. It is interesting to note the existence of small negative

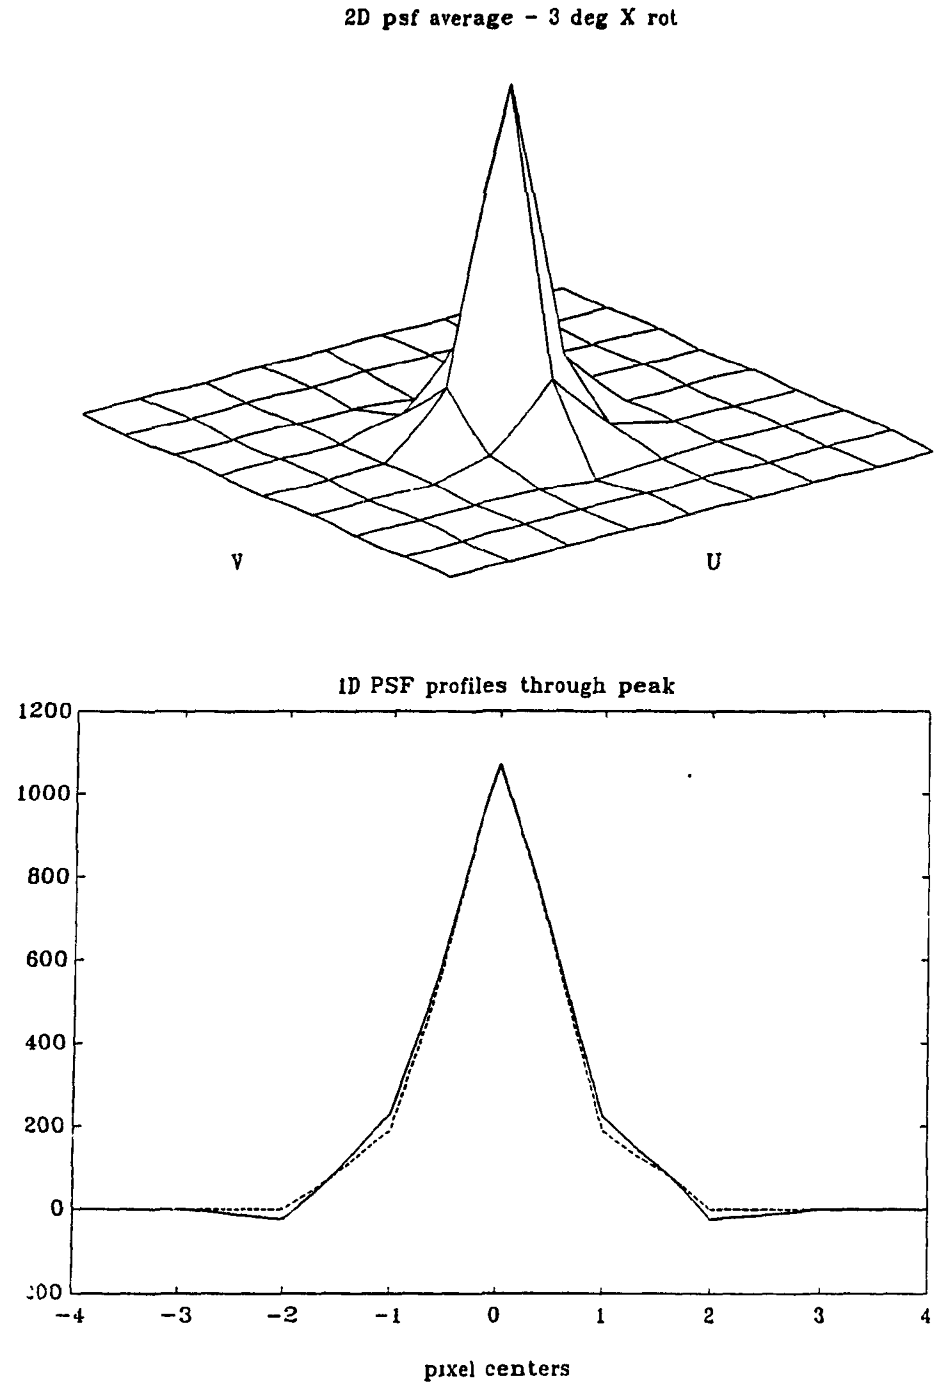

Figure 5.3: PSF of voxel after translations and rotation about x-axis. a) shows 2D elevation plot of PSF, b) shows horizontal and vertical profiles through the PSF peak.

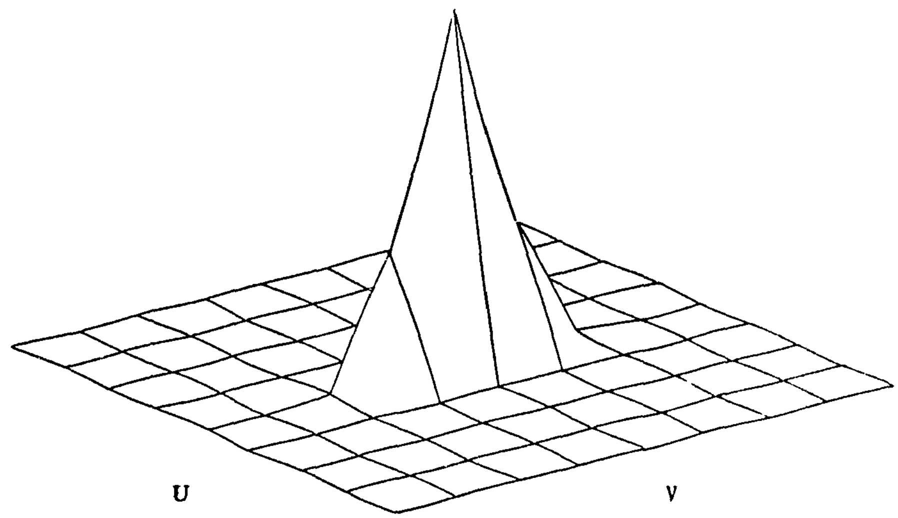

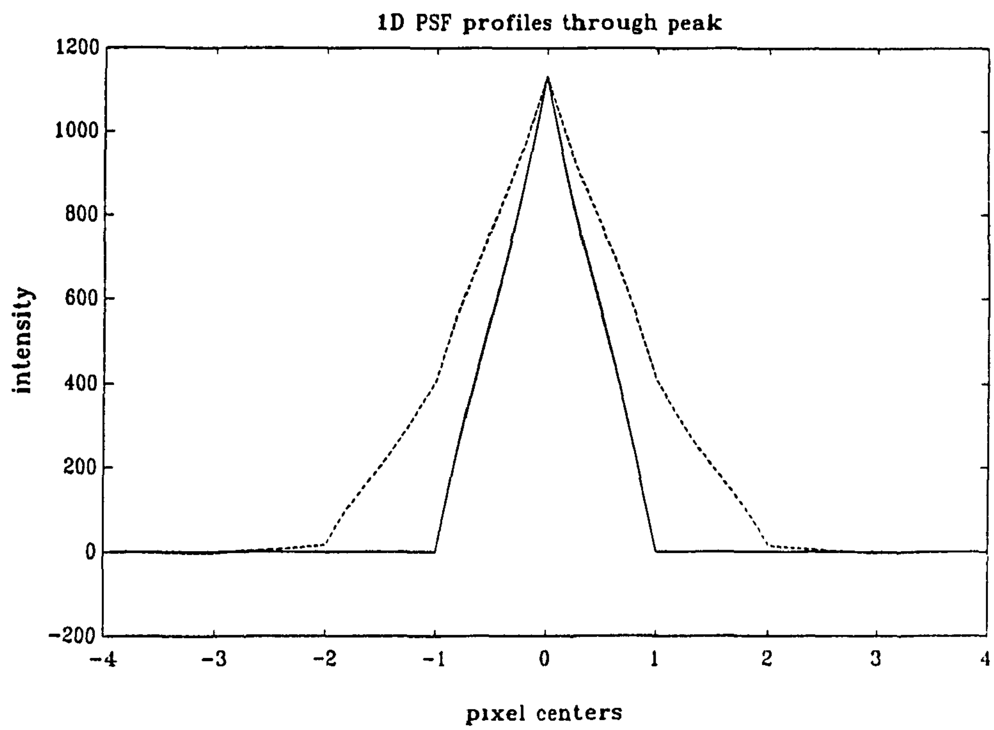

Figure 5.4: PSF of voxel after translations and rotation about y-axis. a) shows 2D elevation plot of PSF, b) shows horizontal and vertical profiles through the PSF peak.

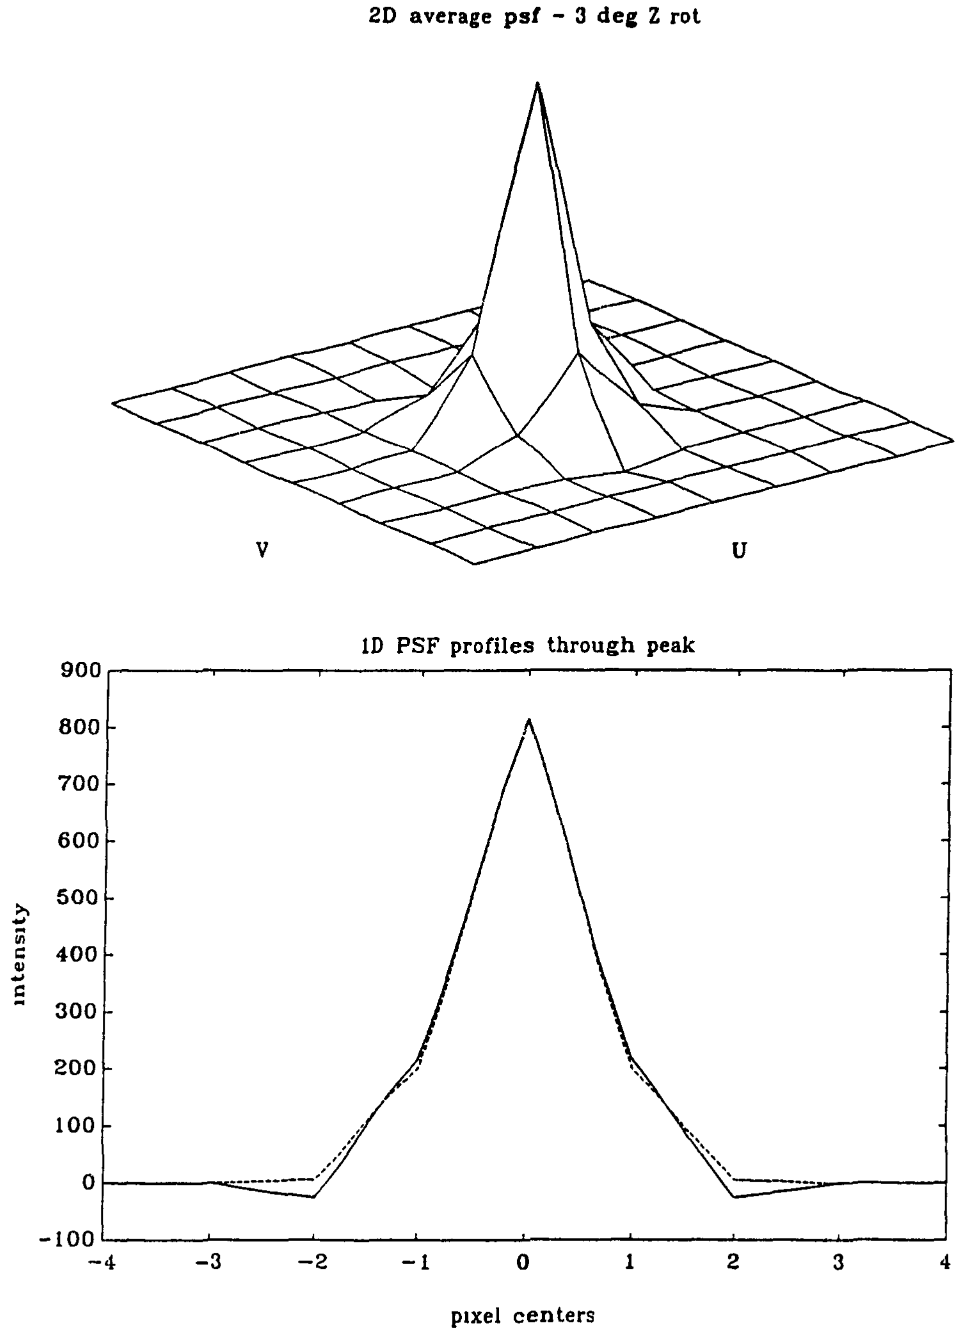

Figure 5.5: PSF of voxel after translations and rotation about z−axis. a) shows 2D elevation plot of PSF, b) shows horizontal and vertical profiles through the PSF peak.

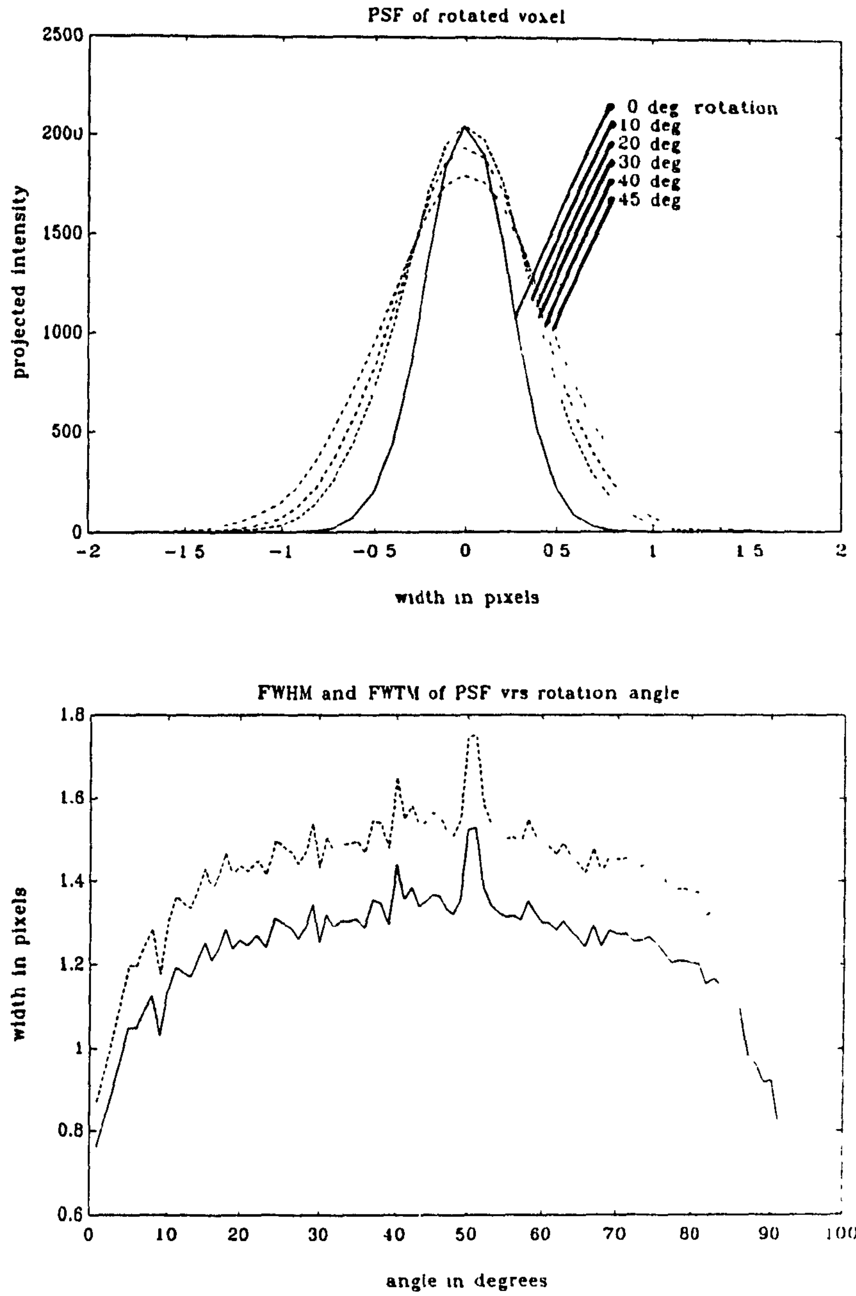

Figure 5.6: PSF of projected voxel vs. rotation angle. a) The graph show the PSF for six different angles of rotation. The PSFs of angles symmetric to 45° are identical. b) The graph show the full-width half-maximum (solid line) and full-width tenth-maximum (dotted line) of the PSF for angles varying from 0° to 90°.

side-lobes in the PSF profiles. This is consistent with the use of a truncated sinc interpolation kernel. The relatively small magnitude of the negative lobes, compared to the size of the central peak is indicative of a gaussian envelope modulating the kernel.

The MTF is calculated directly in the second experiment. A random field is created by generating uniformly distributed random numbers in the range 0 to 511 for each pixel in a 256x256 2D image. The 2D Fourier transform is calculated on a 128x128 window of this input plane using a radix-2 fast Fourier transform (FFT) algorithm [4]. The random field is placed within the volumetric data set in a plane parallel to the projection plane. Translations and rotations are applied to the volume along each axis and the output image is formed using the simple additive projection scheme, as before. The FFT of a 128x128 window of the output image is then calculated. Figure 5.9 shows the plot of the 2D MTF calculated by averaging the ratio of the magnitude of the output matrix FFT over the input FFT for 32 input/output matrix pairs.

Figure 5.9 shows the normalized horizontal and vertical profiles of the MTF through the transform's zero frequency term. The FWHM or 3db point of the PSF is often used to define system resolution, since two points with closer spacing than this will not be resolved. Here, the 3db point on the MTF curve calculated from the PSF is approximately 0.27 cycles per pixel horizontally and 0.31 cycles per pixel vertically. This corresponds to a resolution of 1.6 pixels horizontally and 1.4 pixels vertically. When measured on the MTF of the random field projection, the 3db point falls at 0.32 and 0.34 cycles per pixel horizontally and vertically respectively, corresponding to resolutions of 1.4 and 1.3 pixels in the two directions. The measured value of the FWHM of the rendering PSF is consistent with that theoretically estimated in section 5.4 .

It is interesting to note that the resolution measured from the MTF does not exactly correspond to the FWHM of the PSF in Fig. 5.7. There are two reasons for the discrepency. 1) A single projected point is spread over a small area. Thus, only three to five non-zero samples are available to estimate one of the vertical or horizontal cross-

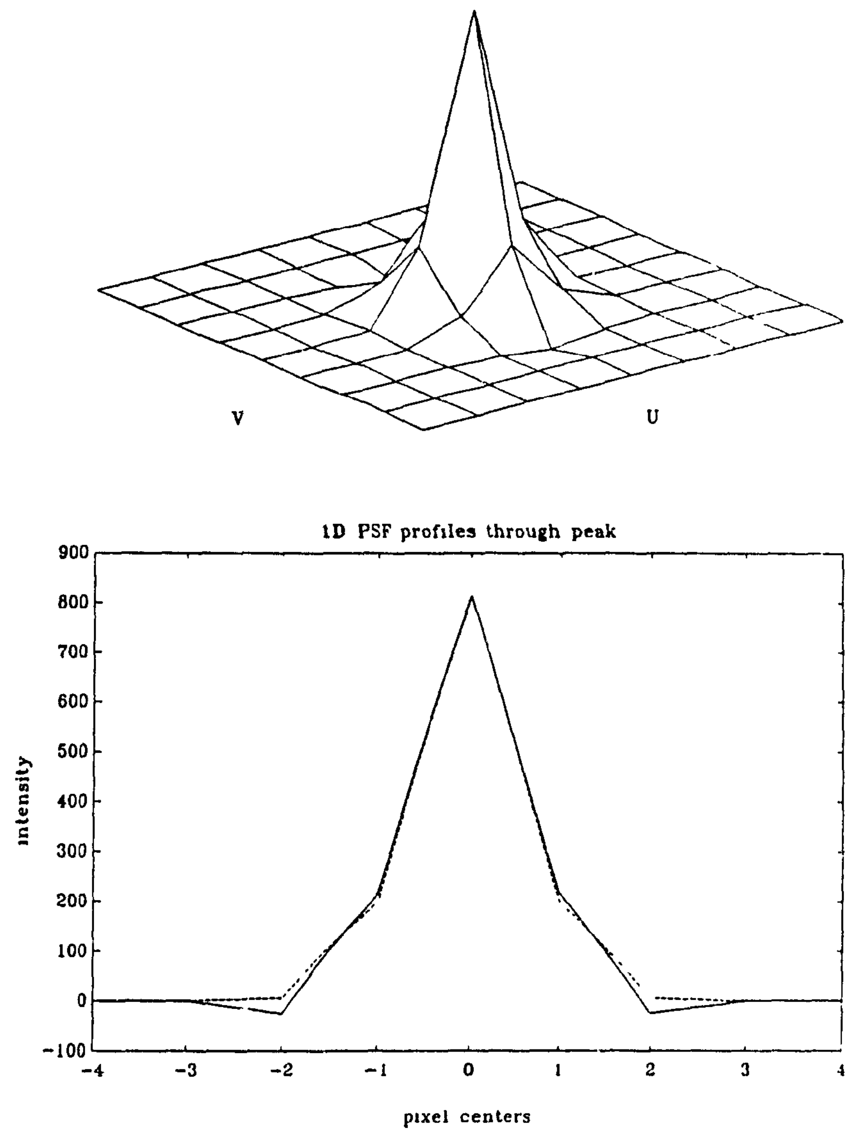

Figure 5.7: PSF of average rendered voxel.

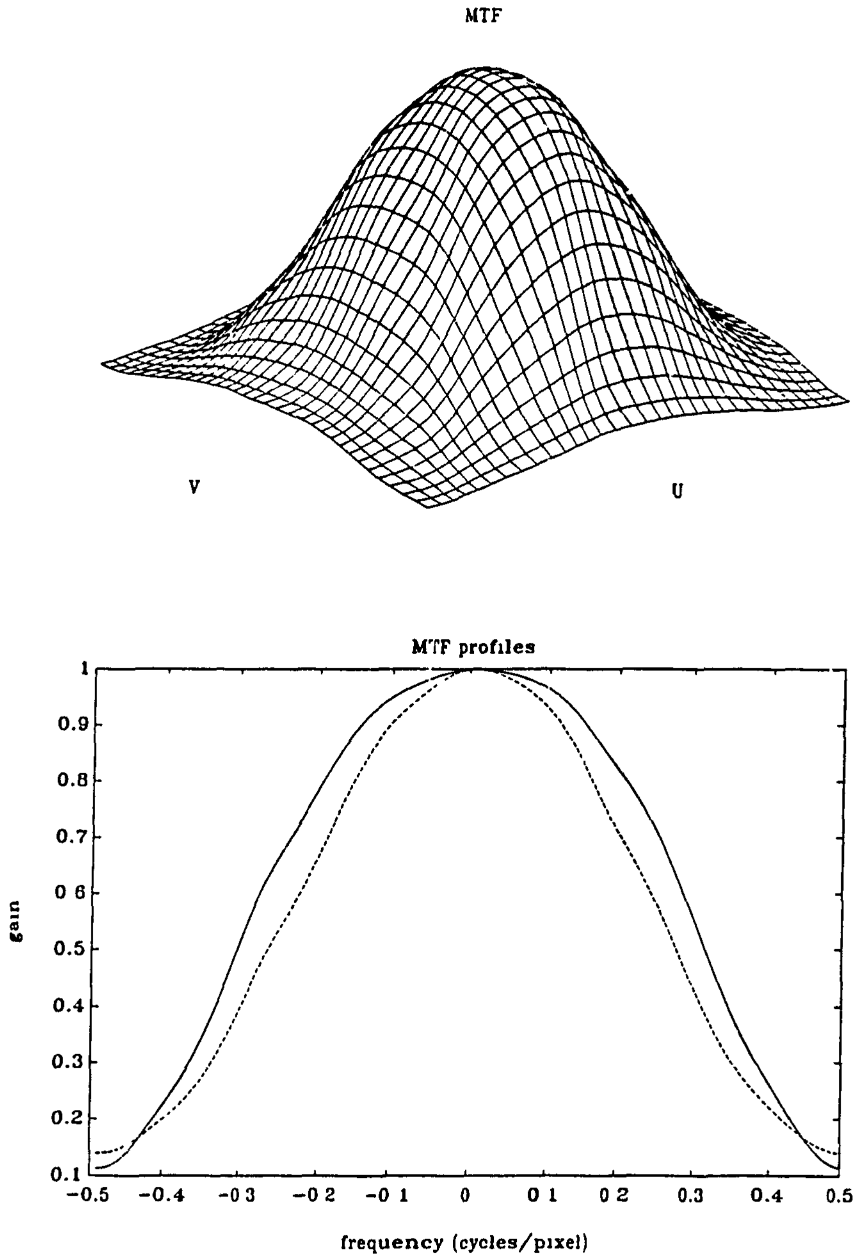

Figure 5.8: 2D MTF of average rendered voxel.

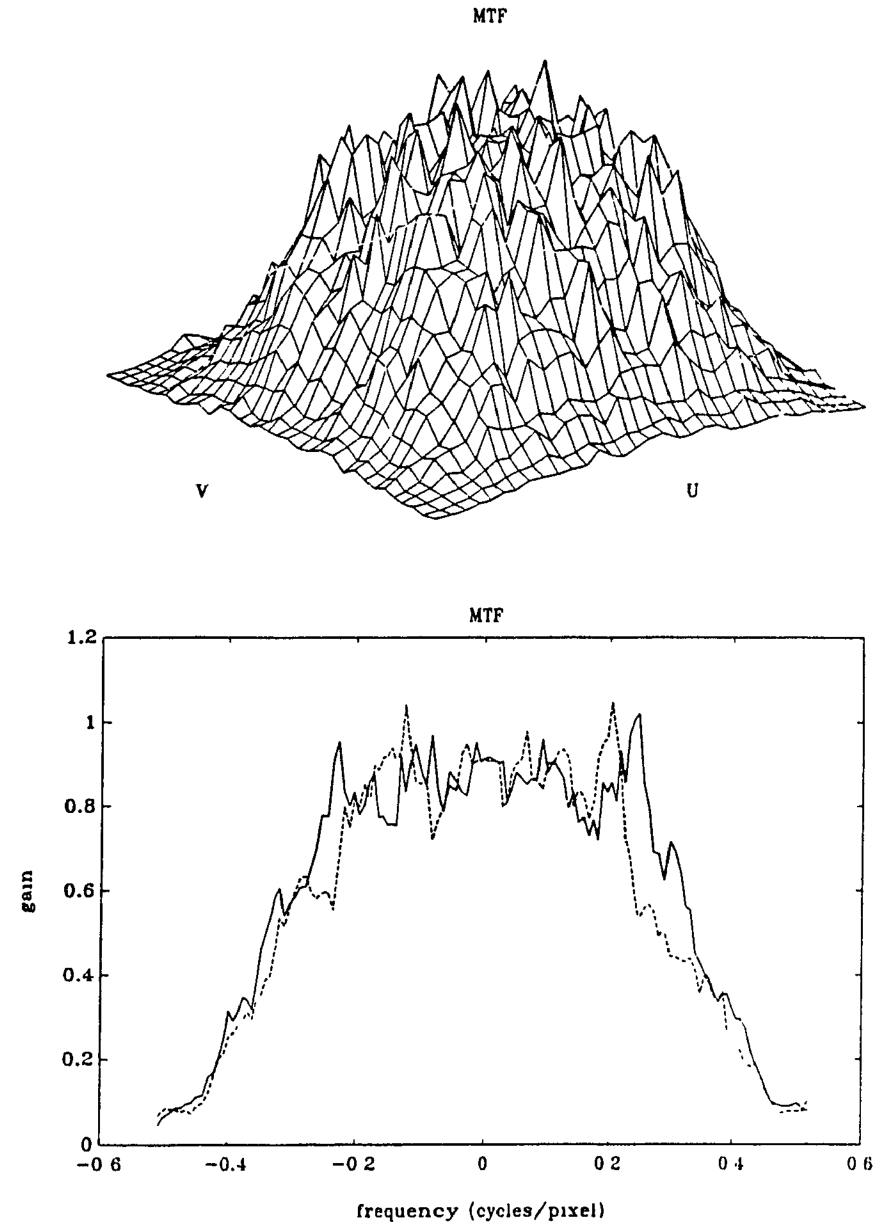

Figure 5.9: 2D MTF calculated from random field in a). Normalized horizontal and vertical profiles in b).

sections of the PSF. This small number of data points limits the refinement of the fitting parameters, even though the experiment was repeated over a large number of points. 2) Experiments have shown that the PSF has the shape of a truncated sinc function. However, a simple gaussian model was used to estimate the FWHM of the PSF.

Other experiments have shown that any 3D point, at a constant distance from the image plane has the same PSF for a given geometric transformation, regardless of its spatial position. The PSF of points at different distances from the image plane have different PSFs due to the differential scaling of the perspective view, however they are directly enlarged or minified by the perspective magnification factor. Since the PSF is point-to-screen distance dependent, the MTF is as well. As the point comes closer to the center of projection, the scaling magnification increases, enlarging the PSF and decreasing the corresponding MTF.

5.6 Summary and discussion

This chapter has presented a derivation of the 2D point spread function of an arbitrary point in a volumetric data set when creating a volume rendered perspective projection in order to quantitatively describe the image quality. The derivation is based on an isotropically sampled band-limited volumetric data set. The processing described in this chapter supposes that the projection operator is of the simple additive type, in order to maintain system linearity.

Experimentation consisting of projecting single voxels and random fields has confirmed the theoretical size estimates of the PSF. An orthogonal projection has a PSF due to translations only, and as shown in Fig. 5.2 is approximately 1.5 pixels FWHM. The PSF of an arbitrary orthogonal projection is between 2.5 to 3.5 pixels FWHM and a perspective perspective view of the volume increases the width of the PSF by the scale of the perspective magnification. The PSF is not necessarily symmetric along each axis, as the width in each direction depends on the projection of the reconstruction kernel in that direction. The measurements of the MTF calculated from the random field confirm these estimates.

The MR tomographic volumes commonly used in medical imaging have a pixel size of 1.2 mm. The results in this chapter indicate that a tomographic voxel is blurred over an approximate extent 4mm in diameter at FWHM in the image plane before perspective magnification. The Siemens Digitron/Angiotron DSA unit has a selectable focal spot size and field of view. When imaging with a focal spot size of 1mm and a field of view of 170 mm, the spatial resolution in the object plane is almost 4 line pairs per mm [35]. A single point is spread to 0.25 mm at FWHM. While there is more than one order of magnitude difference in resolution between the two images, the merged projections are still useful for localization since the tomographic projection image gives an anatomical context for visualization of the vasculature.

Chapter 6

Discussion and Conclusion

6.1 Summary of technique

This thesis has presented methods for merging the complementary information from multiple medical imaging modalities. In particular, the goal of merging tomographic image data (MR and CT) with DSA projections has been realized. The method is based on the presence of fiducial markers during scanning which define each modality's imaging geometry. The perspective projection parameters are calculated from the positions of these fiducial points in the DSA image and then used in the volumetric rendering algorithm to project the tomographic data set onto an equivalant image plane. The two matched projections are then merged by overlaying the DSA image on top of the volume rendering.

Direct volume rendering yields accurate 3D projection images that are free from the artifacts associated with threshold-based rendering methods so that soft, fuzzy or ill-defined surfaces within the volume can be visualized. The rendering technique permits one to directly view 3D anatomical structure previously inferred from 2D tomographic slices. When merged with the DSA projections in stereo, the neurosurgeon is provided with accurate, unclouded 3D images relating anatomy to vasculature.

Since all imaging is performed in the stereotactic environment, the merged images have the same geometry as the DSA projection images. They can be used immediately in the stereotactic planning software to determine the location of target points and position 3D probes, simulating surgery. No extra software tools need to be written to access and analyze these new images.

Viewing the images in stereo permits the analysis to occur in 3D, subject to some constraints [36], which allows the neurosurgeons to consider unorthodox paths for probe placement that would be difficult, if not impossible, to evaluate using a simple 2D biplane technique. The visualization of vessel structure and surrounding tissue in these stereo-pairs allows the neurosurgeon to identify both the target and the vasculature-free path required to attain it without having to refer to multiple images. Selectively choosing regions of the anatomy to be visualized with the vasculature helps to determine their spatial inter-relationships, thus aiding in diagnosis and surgical planning.

Another difficulty in stereotactic procedure planning resides in the selection of a triphenation site[1]. If this hole lies directly above one of the surface vessels on the cortex, it may not be used since any stereotactic instrument entering the hole may pierce the vessel. Volumes of magnetic resonance angiography (MRA) data have been rendered using the modified technique described in chapter 4. The neurosurgeons are excited about the possibility to view, and directly evaluate, the position of vessels on the surface of the cortex. Not only does this allow them to avoid the surface vessels, but major important gyri can be located and circumvented as well.

6.2 Accuracy

Stereotactic localization demands a high degree of measurement accuracy, and an extensive study of this issue has been made with respect to the raw images obtained from each of the modalities [73]. It has been demonstrated that the measured values have an accuracy of ±1mm in DSA, and transverse CT and MRI. 3D image reconstruction

[1]Small burr hole through the skull giving access to a lesion when using stereotactic instruments.

uses many operations that involve interpolation and rotation steps, each of which degrades the modulation transfer function of the imaging operation to some extent. The derivations and experiments in chapter 5 show that the FWHM of the PSF of the rendered images is approximately 4 pixels in diameter, which degrades the resolution of the rendered image to 0.25 line pairs per mm.

6.3 Future work

This thesis has only begun to scratch the surface of 3D medical data visualization through direct volume rendering. Four areas of work may be identified to extend this project: the rendering algorithm itself, preprocessing techniques, other applications and the development of a user interface.

6.3.1 Algorithm

A factor limiting the use of this technique is the algorithm's high computational cost and the resulting lengthly processing time associated with rendering images for each study. The volumetric manipulation and projection for single renderings take approximately 12 to 15 minutes per view and a stereo pair requires 20 to 25 minutes. This is definitely not an interactive technique. However, most of this processing time is due to geometric manipulation of the volume, as only one to two minutes are required for the actual projection. Methods must be examined to determine how the oriented volume can be rendered with different levels of obscuration to aid in the analysis.

While the rendering model is very simple, it should be extended to increase the realism of the rendered images. The lighting model should take into account multiple light sources as well as diffuse and specular reflection from different surface types to create more visible highlights. Shadows could be incorporated into the rendering model as well. (Each of these would incur an additional computational cost.)

6.3.2 Preprocessing

Interpolation of the volumetric data sets in the axial direction is often necessary when dealing with medical tomographic images. The simple interpolation methods used in this thesis result in aliasing artifacts present in the rendered images. A direct solution would be to modify the scanning protocols to obtain more isotropic volumes.

A major obstacle to volume rendering of MR data in this thesis is an automatic, probabilistic, continuous classification scheme. Such a method would permit the use of intelligent interpolation algorithms and the creation of meaningful matte volumes to mask selected regions of data.

6.3.3 Other applications

While merging DSA with CT or MR has been considered and studied in this thesis, techniques to visualize the composition of different volumetric data sets have not been addressed. Interesting results may be obtained using parameters derived from each data set for rendering. For example, the gradient and segmentation information from an MR volume may be used to render PET data in order to visualize the cortical surface activity.

Many institutions render *cine-loops* (short animated sequences) of rotating volumes as an aid to visualization. Cues due to motion give the observer a strong impression of depth. Until now, these sequences have only been used qualitatively. Since the volumetric data sets acquired stereotactically at the MNI contain imaging geometry information, it is possible to project a 3D cursor within the moving volume. The location of the cursor within the image is easily calculated using the volumetric geometry and the projection parameters. As each image in the animation sequence is presented, the cursor position is updated on the screen. This allows cine-loops to be used to obtain quantitative geometric information. This technique will have to be evaluated to determine its use in neurosurgical planning.

Volumetric reformatting (described in chapter 3) is used to examine the data on a slice by slice basis. A ***floating plane*** display of the volumetric data may be achieved by panning through and displaying slices, parallel to the image plane, from two volumes in stereo. The truncated viewing pyramid corresponding to each eye's view would be used as the volumetric dataset. The planes could be merged directly with the DSA projections, yielding yet anouther method to merge projection images with tomographic data.

6.3.4 Interface

A remaining issue is the design of an effective user interface to the system. In clinical practice, the surgeon must be able to quickly retrieve the required data sets from the relevant imaging systems, interact readily with the images in 3D using simple spatial pointing systems, refer any operation back to the original 2D images for confirmation purposes, and to quickly manipulate the 3D image (generation of alternate views, interactive 'dissection' of the images, etc...). Techniques must also be developed to assist with, or to automatically find, suitable rendering parameters for a given data set.

6.4 Synopsis

The technique of merging DSA images with tomographic data provides another view of existing data so that more informed, and thus better, safer decisions regarding surgery planning can be made. Radiologists and neurosurgeons at the MNI are enthusiastic about the work and are looking forward to using it in everyday practice. The initial experience with the system, along with this enthusiasm, leads one to believe that such a system will be a valued addition to the tools used by the neurosurgical community.

Glossary

- **CHAP:** CHannel Array Processor. The PIXAR uses a four channel array processors to quickly manipulate pixel window data. Each processor is applied to one channel of the pixel data.

- **cropping:** defines the sub-region of the projected image to be viewed. Cropping is achieved by translating the origin of the image plane.

- **CT:**Computed Tomography.

- **direct volume rendering:** a technique used to create a 3D view of a volumetric data set by projecting each voxel directly onto the image plane.

- **DSA:** Digital Subtraction Angiography.

- **focal distance:** the distance between the X-ray source and the film plane in the DSA projection. This distance is equal to the perpendicular distance between the center of the projection and the image plane in the rendered image.

- **homogeneous transformation matrix:** a 4×4 matrix which determines how a point (X, Y, Z) in the frame coordinate system is mapped onto the image plane.

- **image plane:** the planar surface (perpendicular to the principal ray) on which the image of the data volume is projected. Also known as the projection plane.

- **MIPS:** Million Instructions Per Second. A computational speed rating.

- **MR:**Magnetic Resonance

- **orthographic:** a projection employing parallel rays as if the center of projection was located at an infinite focal distance.

- **piercing point:** the point of intersection between the principal ray and the image plane. Also known and the *viewing point* or the *stare point*.

- **perspective:** a projection corresponding to the center of projection located at a finite distance from the image plane. The rays are divergent, emenating from the center of projection.

- **principal ray:** the ray defining the viewing direction, intersecting the image plane perpendicularly.

- **SIMD:** Single Instruction Multiple Data. The architecture of the PIXAR permits 4 streams of data to be processed simultaneous by the same instruction by a four channel array processor.

- **stereotactic frame:** is attached to the patient's cranium. It provides a common frame of reference for scanning and a base for surgical instruments.

- **stereotaxy:** is a field of neurosurgery where small and deep seated cerebral tumours are approached through a small hole in the skull using mechanically guided instruments.

- **volume rendering:** is a technique used to produce a 2D image of a higher dimensional (usually 3D) dataset. The data in the volume is projected onto the image plane either directly, or by fitting geometric primitives to each sample point and projecting the primitive.

- **voxel:** is the smallest 3D element making up a volume. A voxel is the 3D equivalent of a 2D image pixel.

Bibliography

[1] *Pixar Image Computer ChapTools User's Guide.* PIXAR, 2.00 edition, July 1987.

[2] *Pixar Image Computer Programmer's Manual.* PIXAR, 2.00 edition, July 1987.

[3] Ehud Artzy. Display of three-dimensional information in computed tomography. *Computer Vision, Graphics, and Image Processing*, 9:196–198, 1979.

[4] G.D. Bergland. A guided tour of the Fast Fourier Transform. *IEEE Spectrum*, 41–52, July 1969.

[5] Fred L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-11(6):567–585, 1989.

[6] D.A. Bosch. *Stereotactic Techniques In Clinical Neurosurgery.* Springer-Verlag Wien, New York, 1986.

[7] Ronald W. Bracewell. *The fourier Transform and its Applications.* McGraw-Hill, New York, 1978.

[8] R. A. Brown. A computerized tomography-computer graphics approach to stereotaxic localization. *Journal of Neurosurgery*, 50:715–720, 1979.

[9] D.L. Burk, L.A. Cooperstein, G.T. Herman, D.C. Mears, and J.K. Udupa. Acetabular fractures - 3-Dimensional computed tomographic imaging and interactive surgical planning. *CT-Journal of Computed Tomography*, 10(1).1–10, 1986.

[10] Edwin Catmull. Computer display of curved surfaces. In *Proc. IEEE Conf. Comput. Graphics Pattern Recognition Data Struct.*, page 11, May 1975.

[11] L.S. Chen, G.T. Herman, R.A. Reynolds, and J.K. Udupa. Surface shading in the cuberille enviroment. *IEEE Computer Graphics and Applications*, 5:33–43, 1985.

[12] H.N. Christiansen and T.W. Sederberg. Conversion of complex contour line definitions into polygonal element mosaics. *ACM Computer Graphics*, 12:187–192, 1978.

[13] H.E. Cline, W.E. Lorenson, S. Ludke, C.R. Crawford, and B.C. Teeter. Two algorithms for the three-dimensional reconstruction of tomograms. *Medical Physics*, 15(3):320–327, 1988.

[14] Louis Collins, Terry Peters, Sean Marrett, and Chris Henri. Integration of stereoscopic angiography data with volume rendered 3-dimensional CT and MRI images (WIP). In *Radiology 1988 RSNA Scientific Program*, page 369, Chicago, Nov. 27-Dec. 2 1988. Presented at RSNA, Chicago Il.

[15] Larry T. Cook, P. Nong Cook, Kyo Rak Lee, Solomon Batnitsky, Bert Y.S. Wong, Steven L. Fritz, Johnathan Ophir, Samuel J. Dwyer, Lawrence R. Bigongiari, and Arch W. Templeton. An algorithm for volume estimation based on polyhedral estimation. *IEEE Transactions on Biomedical Engineering*, BME-27(9):493–499, 1980.

[16] Robert Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. *ACM Computer Graphics*, 22(4):65–74, 1988.

[17] Robert A. Drebin, Loren Carpenter, and Pat Hanrahan. *Pixar's Volume Rendering Technique*. Technical Report 174, PIXAR, May 1987.

[18] R.O. Duda and P.E. Hart. *Pattern Recognition And Scene Analysis*. Wiley, New York, 1973.

[19] Alan Evans, Sean Marrett, Louis Collins, and Terry Peters. Anatomical-functional correlative analysis of the human brain using three dimensional imaging systems. In *Medical Imaging III*, Newport Beach, California, Jan. 29 - Feb. 3 1989.

[20] Alan C. Evans, C. Beil, Sean Marrett, Chris Thompson, and Antoine Hakkim. Anatomical functional correlation using an adjustable MRI based atlas with PET. *Journal of Cerebral Blood Flow and Metabolism*, 8(4):813–830, 1988.

[21] E.J. Farrell, W.C. Yang, and R.A. Zappulla. Animated 3D CT imaging. *IEEE Computer Graphics and Applications*, 5(12):26–32, 1985.

[22] E.K. Fishman, A.F. Brooker, D. Magid, and S.S. Siegelman. Fractures of the sacrum and sacroiliac joint - evaluation by computerized-tomography with multiplanar reconstruction. *Southern Medical Journal*, 81(2):171–177, 1988.

[23] J.D. Foley and A. van Dam. *Fundamentals of Interactive Computer Graphics*. Addison–Wesley, Reading, Mass, 1984.

[24] H. Fuchs, Z.M. Kedem, and S.P. Uselton. Optimal surface reconstruction from planar contours. *Communications of the ACM*, 20:693–702, 1977.

[25] Henry Fuchs, Stephen M. Pizer, Jeffery L. Creasy, Jordan B. Renner, and Julian G. Rosenman. Interactive, richly cued shaded display of multiple 3D objects in medical images. In *SPIE: Medical Imaging II*, pages 842–849, 1988.

[26] Mutsuhisa Fujioka, Nagaaki Ohyama, Takao Honda, Junpai Tsujiuchi, Masano Suzuki, Shozo Hashimot, and Shigeto Ikeda. Holography of 3D surface reconstructed CT images. *Journal of Computer Assisted Tomography*, 12(1):175–178, 1988.

[27] D. Gabor. Holography 1948-1971. *Science*, 177(4046):299–313, 1972.

[28] S. Ganapathy. Decomposition of transformation matrices for robot vision. *IEEE*, 130–139, 1984.

[29] S. M. Goldwasser and R. Anthony Reynolds. Real-time display and manipulation of 3D medical objects: the voxel processor architecture. *Computer Vision, Graphics, and Image Processing*, 39:1–27, 1987.

[30] Rafael C. Gonzalez and Paul Wintz. *Digital Image Processing.* Addison-Wesley, Reading, Mass., 1977.

[31] Dan Gordon and R. Anthony Reynolds. Image space shading of three-dimensional objects. *Computer Vision, Graphics, and Image Processing*, 361–376, 1985.

[32] H. Gouraud. Computer display of curved surfaces. *Transactions of the IEEE*, C-20:623-628, 1971.

[33] L.D. Harris, J.J. Camp, E.L. Ritman, and R.A. Robb. 3-Dimensional display and analysis of tomographic volume images utilizing a varifocal mirror. *IEEE Transactions on Medical Imaging*, 5(2):67–72, 1986.

[34] William R. Hendee. Variables which affect image clarity. In Arthur G. Haus, editor, *The physics of medical imaging: Recording system measurements and techniques*, pages 427–441, American Institue of Physics, Inc., 1979.

[35] Chris Henri. *Application of stereoscopic digital subtraction angiograms to stereotactic neurosurgery planning.* Master's thesis, McGill University, Montreal, 1989.

[36] Chris Henri, Louis Collins, Terry Peters, Alan Evans, and Sean Marrett. Three-dimensional interactive display of medical images for stereotactic neurosurgery planning. In *Medical Imaging III*, Newport Beach, California, Jan. 29 - Feb. 3 1989.

[37] Gabor T. Herman. *Image Reconstruction from Projections: the fundamentals of computerized tomography.* Academic Press, New York, 1980.

[38] G.T. Herman and H.K. Liu. Three-dimensional display of human organs from computed tomograms. *Computer Graphics and Image Processing*, 9:1–21, 1979.

[39] Karl Heinz Höhne and Ralph Bernstein. Shading 3D-images from CT using gray-level gradients. *IEEE Transactions on Medical Imaging*, MI-5(1):45–47, 1986.

[40] Karl Heinz Höhne, Robert L. Delapaz, Ralph Bernstein, and Robert C. Taylor. Combined surface display and reformatting for three-dimensional analysis of tomographic data. *Investigative Radiology*, 22:658–664, 1987.

[41] X. Hu, D.N. Levin, S.G. Galhorta, C.A. Pelizzari, G.T.Y. Chen, R.N. Beck, C-T. Chen, and M.D. Cooper. Volumetric rendering of multimodality, multivariable medical image data. In *Chapel Hill Workshop on Volume Visualization*, pages 45–50, Chapel Hill, NC, May 18-19 1989.

[42] T.S. Huang. Digital holography. *Proceedings of the IEEE*, 1335–1346, Sept. 1971.

[43] I.T. Jackson and U. Bite. 3-Dimensional computed tomographic scanning and major surgical reconstruction of the head and neck. *Mayo Clinic Proceedings*, 61(7):546–555, 1986.

[44] E Keppel. Approximating complex surfaces by triangulation of contour lines. *I.B.M. Journal of Research and Development*, 19:2–11, 1975.

[45] E.S. Kerekes. A simple devise for stereoscopic viewing of films. *American Journal of Roentgenology*, 75(1):140, 1956.

[46] U. Kretschmar and S. Schneider. Intra-arterial DSA in routine diagnostic neuroradiology. *Medicamundi*, 33(3):117–124, 1988.

[47] Anil Kumar, Dieter Welti, and Richard R. Ernst. NMR Fourier Zeugmatography. *Journal of Magnetic Resonance*, 18:69–83, 1975.

[48] J.M. Lane, L.C. Carpenter, T. Whitted, and Blinn J.F. Scan line methods for displaying parametrically defined surfaces. *Communications of the ACM*, 23(23):23–34, 1980.

[49] L. Leksell. *Stereotaxis And Radiosurgery: An Operative System.* Charles C. Thomas, Springfield, Illinois, 1971.

[50] L. Leksell, C. Lindquist, J. Adler, D. Leksell, B. Jernberg, and L. Steiner. A new fixation device for the Leksell stereotactic system. *Applied Neurophysiology*, 50:9–22, 1987.

[51] David N. Levin, Charles A. Pelizzari, George T.Y. Chen, Chin-Tu Chen, and Malcolm D. Cooper. Retrospective geometric correlation of MR, CT and PET images. *Radiology*, 169:817–823, 1988.

[52] Marc Levoy. Design for a real-time high quality volume rendering workstation. In *Chapel Hill Workshop on Volume Visualization*, pages 85–92, Chapel Hill, NC, May 18-19 1989.

[53] Marc Levoy. Direct visualization of surfaces from computed tomography data. In *SPIE: Medical Imaging II*, pages 828–841, 1988.

[54] Mark Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 29–37, May 1988.

[55] Cheng-Chung Liang, Wei-Chung Lin, and Chin-Tu Chen. Intensity interpolation for serial cross-sections. In *SPIE: Medical Imaging III*, pages 60–63, 1989.

[56] W.E. Lorenson and H.E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. *ACM Computer Graphics*, 21:163–169, 1987.

[57] L. D. Lundsford, A. J. Martinez, and R. E. Latchaw. Stereotaxic surgery with a magnetic resonance and computerized tomography-compatible system. *Journal of Neurosurgery*, 64:872–878, 1986.

[58] L.D. Lunsford, editor. *Modern Stereotactic Neurosurgery.* Martinus Nijhoff Publishing, Boston, 1988.

[59] D. Magid, A.F. Brooker, E.K. Fishman, B.R. Mandelbaum, and S.S. Siegelman. Multiplanar computed-tomography of acetabular fractures. *Journal of Computer Assisted Tomography*, 10(5):778–783, 1986.

[60] B.R. Mandelbaum, M. Bosse, A.F. Brooker, A.R. Burgess, and D. Fishman, E.K. Magid. Multiplanar computed-tomography - a multidimensional tool for evaluation and treatment of acetabular fractures. *Journal of Computer Assisted Tomography*, 11(2):167–173, 1987.

[61] J Mazziota and H.K. Huang. THREAD (three dimensional reconstruction and display) with biomedical applications in neuron ultra-structure and computerised tomography. *American Federation of Information Processing Societies*, 241–250, 1976.

[62] Donald J. Meagher. Efficient synthetic image generation of arbitrary 3D objects. In *Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing*, pages 473–478, June 1982.

[63] William M. Newman and Robert F. Sproull. *Principles of Interactive Computer Graphics.* McGraw–Hill, New York, 1979.

[64] H. Nyquist. Certain topics in telegraph transmission theory. *A.I.E.E. Trans.*, 617, 1928.

[65] A. Olivier, G. Bertrand, and T. Peters. Stereotactic systems and procedures for depth electrode placement: technical aspects. *Applied Neurophysiology*, 46:37–40, 1983.

[66] Andre Olivier and Gilles Bertrand. *Stereotaxic implantation of depth electrodes for seizure recording.* Technical Report, Montreal Neurological Institute and Hospital, 1982.

[67] Andre Olivier, Terry Peters, and Guy Bertrand. Stereotactic system and apparatus for use with MRI, CT and DSA. *Applied Neurophysiology*, 48:94–96, 1985.

[68] Alan V. Oppenheim and Ronald W. Schafer. *Digital Signal Processing*. Prentice Hall, Englewood Cliffs, N.J., 1975.

[69] C.A. Pelizzari, G.T.Y. Chen, D.R. Spelbring, R.R Weichselbaum, and C.T. Chen. Accurate three-dimensional registration of CT, PET and MR images of the brain. *submitted to JCAT*, 1988.

[70] T. M. Peters. Principles and applications of magnetic resonance imaging MRI in neurology and neurosurgery. *The Journal of Mind and Behavior*, 9(3):241–262, 1988.

[71] T. M. Peters and B. C. Sanctuary. *Overview of NMR Reconstruction Principles*, pages 15–26. Society of Nuclear Medicine, 1984.

[72] Terry Peters, John Clark, Andre Olivier, Erich Marchand, George Mawko, Martial Dieumegarde, Letitia Muresan, and Romeo Ethier. Integrated stereotaxic imaging with CT, MR imaging and Digital Subtraction Angiography. *Radiology*, 161:821–826, 1986.

[73] Terry Peters, John Clark, Bruce Pike, Louis Collins, Chris Henri, D. Leksell, and O. Jeppsson. Stereotactic neurosurgery planning on a PC-based workstation. In *SPIE Medical Imaging III*, Newport Beach, California, Jan. 29 - Feb. 3 1989.

[74] Terry M. Peters, John Clark, Bruce Pike, M. Drangova, and A. Olivier. Stereotactic surgical planning with magnetic resonance imaging, digital subtraction angiography and computed tomography. *Applied Neurophysiology*, 50:33–38, 1987.

[75] Terry M. Peters, A. Olivier, and G. Bertrand. The role of computed tomographic and digital radiographic techniques in stereotactic procedures for electrode implantation and mapping, and lesion localization. *Applied Neurophysiology*, 46:200–205, 1983.

[76] Bui-Tong Phong. Illumination for computer generated images. *Communications of the ACM*, 18:311–317, 1975.

[77] Bruce Pike, Ervin B. Podgorsak, Terry M. Peters, and Conrado Pla. Dose distributions in dynamic stereotactic radiosurgery. *Medical Physics*, 14(5):780–789, 1987.

[78] G. Bruce Pike, Terry M. Peters, Edwin B. Podgorsak, Conrado Pla, A. Olivier, and A. de Lotbinière. Stereotactic external beam calculations for radiosurgical treatment of brain lesions. *Applied Neurophysiology*, 50:269–273, 1987.

[79] G.B. Pike, L. Collins, T.M. Peters, R. del Carpio, R. Whitman, and S. McLachlan. Acquisition and display of 3D Magnetic Resonance Angiograms. *sumitted to Radiology*, 1989.

[80] Thomas Porter and Tom Duff. Compositing digital images. *ACM Computer Graphics*, 18(3):253–259, 1984.

[81] William K. Pratt. *Digital Image Processing.* John Wiley and Sons, New York, 1978.

[82] Ian Pykett, Jeffrey Newhouse, Ferdinando Buonanno, Thomas Brady, Mark Goldman, Phillip Kistler, and Gerald Pohost. Principles of nuclear magnetic resonance imaging. *Radiology*, 143:157–168, April 1982.

[83] Osman Ratib, Luc Bidaut, Heinrich R. Schelbert, and Michael E. Phelps. A new technique for elastic registration of tomographic images. In *SPIE: Medical Imaging II*, pages 452-455, 1988.

[84] R. Anthony Reynolds, Dan Gordon, and Lih-Shyang Chen. A dynamic screen technique for shaded graphics display of slice-represented objects. *Computer Vision, Graphics, and Image Processing*, 38:275-298, 1987.

[85] R.A. Reynolds, L.S. Chen, and M.R. Sontag. An algorithm for 3-Dimensional visualization of radiation- therapy beams. *Medical Physics*, 15(1):24-28, 1988.

[86] R.A. Robb, J.J. Camp, D.P. Hanson, and P.B. Heffernan. A workstation for multi-dimensional display and analysis of biomedical images. *Computer Methods and Programs in Biomedicine*, 25(2):169–184, 1987.

[87] D.D. Robertson, E.K. Fishman, J.W. Granholm, D. Magid, P.S. Nelson, P.C. Walker, and P.J. Weiss. Design of custom hip stem prostheses using 3-Dimensional CT modeling. *Journal of Computer Assisted Tomography*, 11(5):804–809, 1987.

[88] David F. Rogers. *Procedural Elements for Computer Graphics.* McGraw-Hill, New York, 1985.

[89] G. Russell and R.B. Miles. Display and perception of 3-D space-filling data. *Applied Optics*, 26(6):973–982, 1987.

[90] Alvy Ray Smith. *Volume Graphics and Volume Visualization, a tutorial.* Technical Report 176, PIXAR, May 1987.

[91] R.L. Stern, H.E. Cline, G.A. Johnson, and C.E. Ravin. Development of a 3D reconstructed-image surgical planning station. *Investigative Radiology*, 23(9):60, 1988.

[92] T.M. Strat. Recovering the camera parameters from a transformation matrix. In *Proc. DARPA IU Workshop*, pages 265–271, Oct. 2-3 1984.

[93] P. Suetens and A. Oosterlinck. Pseudo-holographic display system for integrated 3-D medical images. In *SPIE: Medical Imaging*, pages 606–613, 1987.

[94] I.E. Sutherland. Three-dimensional data input by tablet. *Proceedings of the IEEE*, 62:453–471, 1974.

[95] M. Takahashi, Y. Ozawa, and H. Takemoto. Focal-spot separation in stereoscopic magnification radiography. *Radiology*, 140:227–229, 1981.

[96] B.A. Toth, D.S. Ellis, and W.B. Stewart. Computer-designed prostheses for orbitocranial reconstruction. *Plastic and Reconstructive Surgery*, 81(3):315–322, 1988.

[97] Heang K. Tuy and Lee Tan Tuy. Direct 2-D display of 3-D objects. *IEEE Computer Graphics and Applications*, 4(10):29–33, October 1984.

[98] J.K. Udupa. Interactive segmentation and boundary surface formation for 3-D digital images. *Computer Graphics and Image Processing*, 18:213–235, 1982.

[99] D.J. Valentino, J.C. Mazziotta, and H.K. Huang. Mapping brain function to brain anatomy. In *SPIE: Medical Imaging II*, pages 445–450, 1988.

[100] Michael W. Vannier, Jeffery L. Marsh, and James O. Warren. Three dimensional computer graphics for craniofacial surgical planning and evaluation. *ACM Computer Graphics*, 17(3):263–273, 1983.

[101] Lee Westover. Interactive volume rendering. In *Chapel Hill Workshop on Volume Visualization*, pages 9–18, Chapel Hill, NC, May 18-19 1989.

[102] C. Worthington, T.M. Peters, R. Ethier, D. Melanson, J. Theron, J.-G. Villemure, A. Olivier, J. Clark, and G. Mawko. Stereoscopic digital subtraction angiography in neurological assessment. *American Journal of Neuroradiology*, 6:802–808, 1985.

[103] T. Wright. A one-pass hidden-line remover for computer drawn three-space objects. In *Proc. 1972 summer Computer Simulation Conference*, pages 261–267, 1972.

[104] S.B. Xu and W.X. Lu. Surface reconstruction of 3D objects in computerized-tomography. *Computer Vision, Graphics, and Image Processing*, 44(3):270–278, 1988.