

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

**Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA**

UMI[®]
800-521-0600

Computer guidance for thalamotomy and pallidotomy

by

Philippe St-Jean

Medical Physics Unit
McGill University, Montréal

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

©P. St-Jean, September 1997.



National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-44289-6

Canada

Contents

List of Figures	vi
Acknowledgments	vii
Abstract	ix
Résumé	x
Contributions	xi
List of Symbols and Abbreviations	xii
Preface	xiv
Introduction	xvi
1 Stereotaxy and Image-Guided NeuroSurgery	1
1.1 Conventional Stereotaxy	1
1.2 Stereotactic Frame	3
1.3 Image-Guided Neurosurgery	6
2 Thalamotomy and Pallidotomy	10

2.1	Symptomatic Parkinson's syndrome	10
2.2	History and description of thalamotomy and pallidotomy	11
2.3	Conventional imaging modalities for thalamotomy and pallidotomy	14
3	Atlases and Automatic Labeling of the Deep Brain Structures	17
3.1	Schaltenbrand and Bailey Atlas	18
3.2	Model MRI Data Set: Superbrain	24
3.3	Patient space	27
3.4	On the Different Spaces and Coordinate Systems	28
3.4.1	Atlas Space (or Schaltenbrand Space):	29
3.4.2	Model MRI Space (or Superbrain Space):	29
3.4.3	Native MRI space (Pre-op, Intra-op, Post-op):	29
3.4.4	Frame Space:	29
3.4.5	Physical Space (Brain Space, Patient Space):	29
4	Image-Guided NeuroSurgery Tools for Thalamotomy and Pallidotomy	31
4.1	Neurological Navigation in VIPER	32
4.1.1	Trajectory Views	37
4.1.2	Markers	38
4.2	3D Tools in VIPER	39
4.2.1	Surface Rendering	39
4.2.2	Volume Rendering	45
4.3	Neurosurgical Tools	48
4.4	Stereoscopic Viewing	53

5	Lesion Modeling	56
5.1	Rationale	56
5.2	Algorithm	58
5.3	Virtual Lesion Generator	61
6	Clinical Considerations	65
6.1	Step-By-Step Clinical Procedure	65
6.1.1	Pre-operative steps	65
6.1.2	After the frame is attached to the patient	65
6.1.3	During the operation	66
6.1.4	Post-operative steps	67
6.2	Clinical Results	67
7	Conclusion and Future Work	69
7.1	Conclusion	69
7.2	Future Work	70
	Bibliography	71

List of Figures

1.1	OBT (Olivier-Bertrand-Tipal) stereotactic frame with the fiducial marker plates. . . .	4
1.2	Arc system mounted on the stereotactic frame.	5
1.3	The FARO Surgical Arm.	7
1.4	The Viewing Wand IGNS system main screen.	8
2.1	The Vim node of the thalamus and the Gpm.	12
2.2	Optical scans of the leukotome and the stimulator with different protusions of their wires	14
2.3	Ventriculograms.	15
3.1	Schaltenbrand atlas.	19
3.2	Construction of the volumetric atlas.	21
3.3	Two views of the stacked slices.	22
3.4	Our MRI model data set: Superbrain	26
3.5	Part of the tagging performed between the Atlas (left) and the MRI model (right). . .	27
4.1	The VIPER platform.	33
4.2	Coronal view of a patient's MRI with PET and MRSI scans superimposed	36
4.3	The two trajectory view modes in VIPER.	38
4.4	Surface rendering of the Superbrain's ventricles in VIPER.	42

4.5	White-matter weighted volume rendering of the Superbrain.	47
4.6	Grey-matter weighted volume rendering of the Superbrain.	47
4.7	Example of the orthogonal projection ambiguity.	49
4.8	Standard tri-planar views and trajectory view of the leukotome model, the atlas and an MRI data set.	52
4.9	The 3D leukotome model inside the Vim node.	53
4.10	Stereoscopic image pair showing a mesh model of the cortex arranged for cross-eyed viewing.	54
5.1	The leukotome loop intersecting two different planes at two different distances from the leukotome shaft.	59
5.2	The boundary-finding algorithm in one plane.	60
5.3	A virtual lesion (white) superimposed on the atlas in 2D viewing mode.	62
5.4	A 3D view of a virtual lesion (flat grey) of the Gpm (textured grey).	63
5.5	The Lesion Modeling graphical interface (left).	64
6.1	2D views of a post-operative MRI with superimposition of the atlas.	67

Acknowledgments

I would like to thank:

- Dr Abbas Sadikot for his guidance, motivation and support in relation to all aspects of the project,
- Diego Clonda for his collaboration during the creation of the first version of the atlas, for providing the second atlas version and for great help during preparation for clinical cases,
- Reza Kasrai for his work on the VIPER platform, and for answering my numerous questions about the VIPER coding structure,
- Roch Comeau for his advice in computer graphics, for some VIPER code and for all manner of technical help,
- Rob Flanz for providing the basic structure to VIPER,
- Patrice Munger and Dr Bruce Pike for their collaboration on MRI-related issues,
- Martin Cyr for his help in the clinical environment,
- Mike Sinasac for further VIPER developments,
- Dr Noor Kabani for the manual tagging of the atlas to the model MRI,
- Dr Louis Collins and David McDonald for the ANIMAL algorithm and the Display program respectively,

- The people of the Neuro-Imaging Lab and the McConnell Brain Imaging Centre for general technical support,
- Anne-Marie DeBlois, Maurice St-Jean and Christiane St-Jean for their love and patience.

Finally, I want to thank my supervisor, Dr Terry Peters, for his extraordinary devotion and constant support for my project, for his useful advise and guidance, and for reading and suggesting corrections to the manuscript of this thesis on such short notice.

Abstract

Thalamotomy and pallidotomy constitute the two principal neurosurgical treatments for Parkinson's disease. Both the thalamus nuclei and the internal globus pallidus structures are found very close to the internal capsule, a critical structure that must be avoided in performing the lesion. This forces the lesion of parts of the thalamus and the globus pallidus medialis to be achieved with a spatial precision on the order of 2mm. The leukotome is a neurosurgical tool consisting of a 20cm rigid shaft with a thin metallic wire attached to its tip that can be taken in and out of the body of the shaft. It is used as a knife to perform the excision.

Some critical basal ganglia structures show no anatomical differences at all on MRI scans, even though they have different functionality. In order to provide the neurosurgeon with this missing information, we have developed a deformable volumetric atlas of the basal ganglia from the cryogenic slices atlas of Shaltenbrand and Wahren. The atlas can be non-linearly deformed to fit a patient's MRI.

We have developed a visualization platform for pre-operative and intra-operative utilization of the atlas and of the MRI datasets. Any number of datasets can be superimposed (MRI, atlases, virtual lesions...). Merging of datasets with different extents and different resolutions is supported. The platform includes 3D visualization tools as well. Graphical tools allow the neurosurgeon to see projections of the leukotome from any point of view over the MRI data and the atlas. The software is able to create models of lesions that can be compared with the MRI data and the atlas. It can also suggest an operation protocol (how to use the leukotome) if a target volume is given. This target volume can easily be drawn on the atlas by the neurosurgeon.

Résumé

La thalamotomie et la pallidotomie constituent les deux principaux traitements neuro-chirurgicaux à la maladie de Parkinson. Les noyaux thalamiques ainsi que le globus pallidum mediale impliqués dans ces ablations sont tous situés près de la capsule interne, une structure critique qu'il faut éviter lors de l'ablation. Cette restriction impose que la région à lésionner soit identifiable à une précision spatiale de l'ordre de 2mm. L'ablation est opérée à l'aide d'un leukotome, un outil neuro-chirurgical constitué d'une tige rigide de 20 cm et d'un mince fil métallique attaché à son extrémité et pouvant sortir de la tige pour former une boucle. Cette boucle est utilisée comme lame afin de créer une lésion dans la zone cible.

Certaines structures critiques de la région thalamique ne démontrent aucun contraste sur les images de résonance magnétique (IRM), malgré le fait qu'elles accomplissent des fonctions différentes. Nous avons construit un atlas volumétrique déformable des régions thalamiques à partir de tranches cryogéniques provenant de l'atlas de Schaltenbrand et Wahren. Cet atlas peut être déformé non-linéairement pour correspondre aux IRM de n'importe quel patient.

Nous avons créé une plateforme de visualisation destinée à une utilisation pré- et intra-opératoire de l'atlas et des IRM. Un nombre arbitraire de données visuelles peuvent y être fusionnées et visualisées (IRM, atlas, lésions virtuelles). Des images de tailles et de résolutions différentes peuvent être fusionnées. La plateforme propose aussi des outils de visualisation tridimensionnelle. Elle permet également au neuro-chirurgien de visualiser le leukotome combiné à l'atlas et à l'IRM. Le programme peut créer des lésions virtuelles pouvant être comparées à l'atlas et à l'IRM. Il peut aussi suggérer un protocole opératoire (sur l'utilisation du leukotome) à partir d'un volume-cible, qui peut être déterminé à même l'atlas par le neuro-chirurgien.

Contributions

Philippe St-Jean, D. Clonda, A. Sadikot, A. Evans and T. Peters. Computer guidance for thalamotomy and pallidotomy. In *Proceedings of the Canadian Organisation of Medical Physicists*, pages 146-148, 1997.

Diego Clonda, P. St-Jean, A. Sadikot, T.M. Peters, A.C. Evans. Automatic labeling of the patient's brain with a subcortical atlas in image guided neurosurgery. In *Proceedings of the Canadian Organisation of Medical Physicists*, pages 209-211, 1997.

R. Kasrai, P. St-Jean, M.C. Preul, S. Narayanan, D.L. Arnold and T.M. Peters. Multi-modality image visualization for image-guided neurosurgery. In *Proceedings of the Canadian Organisation of Medical Physicists*, pages 137-139, 1997.

R. Kasrai, P. St-Jean, D. Clonda, A. F. Sadikot, A. C. Evans, and T. M. Peters. Interactive 3-Dimensional Visualization Tools for Stereotactic Atlas-Based. Functional Neurosurgery. Submitted to *Proceedings of SPIE (Medical Imaging) 1998. San Diego, CA.*

List of Symbols and Abbreviations

2D	Two-Dimensional
3D	Three-Dimensional
AC	Anterior Commissure
ANIMAL	Automatic Nonlinear Image Matching and Anatomical Labelling
C, C++	C and C++ computer languages
CPU	Co-Processor Unit
CSF	Cerebro-Spinal Fluid
CT	Computerized Tomography
DSA	Digital Subtraction Angiography
EEG	Electro-Encephalography
fMRI	functional Magnetic Resonance Imaging
Gpl	Globus pallidus laterale
Gpm	Globus pallidus mediale
GUI	Graphical User Interface
ICBM	International Consortium on Brain Mapping
IGNS	Image-Guided Neurosurgery
MNI	Montreal Neurological Institute
MRA	Magnetic Resonance Angiography
MRI	Magnetic Resonance Imaging (or Images)
MRSI	Magnetic Resonance Spectroscopic Imaging

OBT	Olivier-Bertrand-Tipal Stereotactic Frame
OR	Operating Room
osc./s	oscillations per second
PC	Posterior Commissure
PET	Positron Emission Tomography
SNR	Signal-to-Noise Ratio
V	Volt
Vci	Ventralis centralis internus nucleus
Vce	Ventralis centralis externus nucleus
Vim	Ventralis intermedius nucleus
VIPER	Visual Integration Platform for Enhanced Reality

Preface

Thalamotomies and pallidotomies have been performed at the Montreal Neurological Institute over the past thirty years. Improvement to the technique included essentially micro-electrode recordings of the basal ganglia electrical activity [4, 39], and computer guidance using a linearly deformable atlas of the thalamic regions and surroundings [6, 3]. These computer guidance tools were developed twenty years ago and have not seen any important advances since.

The advent of 3D computer graphics, along with new powerful visualization techniques, opened the path to a new era in Image-Guided NeuroSurgery. However, in the case of thalamotomy related computer guidance at the MNI, the development of a 3D non-linear deformation algorithm capable of computing a “matching” transformation between two MRI volumes was the real starting point of our work. The ANIMAL algorithm, developed by Dr Louis Collins at the MNI, makes use of a cross-correlation technique to compute the optimal non-linear transformation between two MRI volumetric data sets using a multi-scale approach.

Atlas-based guidance is indicated for thalamotomy and pallidotomy since no imaging modality allows for the precise identification of the small regions involved in this operative procedure. The aim of this project was to utilize the powerful ANIMAL algorithm to provide a non-linearly deformable atlas that can be superimposed on any patient’s MRI data set, for Image-Guided NeuroSurgical use. This atlas idea was proposed by Dr Terry Peters, Dr Alan Evans and Dr Abbas Sadikot.

This extensive project had to be divided in two parts:

- The creation of an easily deformable 3D atlas of the basal ganglia.

- The development of specific IGNS tools for thalamotomy and pallidotomy that allow for the implementation of the atlas-based guidance in the clinical neurosurgical environment.

The first part has been the subject of Diego Clonda's MSc Thesis. The first "rough" version of the atlas was developed by him in close collaboration with the author. The scanning, concatenation and interpolation of the cryogenic stained slices from the Schaltenbrand and Wahren atlas (from which the atlas was constructed) was done by both of us. The second "refined" version of the atlas, which includes many more regions than the first (and which was the one used clinically at the time this thesis was written), is the work of Diego Clonda alone.

The second part of the project is the subject of this thesis. The creation of the VIPER platform, on which most of the thalamotomy-related IGNS tools discussed in the next chapters were developed, is the original work of Rob Flanz and Reza Kasrai from the Neuro-Imaging Lab at the MNI. The implementation of multiple data set visualization (necessary for atlas and MRI visualization) is the fruit of the collaboration between these two and the author. Though this capability of the platform was first developed to allow for atlas visualization, it has found an important application in tumor-removal operations where the superimposition of a large number of PET scans, MRSI scans and a MRI scan has already been proven useful in about thirty cases. It must be noted that an important part of the VIPER code comes from the previous work of Roch Comeau at the MNI.

The 3D visualization features of the VIPER platform were greatly inspired by David MacDonald's work at the MNI. In particular, the creation of the surface rendered object is still done using his DISPLAY software, and the encoding format for these objects in VIPER is the same as his. Implementation of the 3D viewing mode in VIPER is mostly the work of Mike Sinasac and the author, while the stereoscopic visualization feature was implemented by Reza Kasrai. The lesion modeling tools are the sole work of the author.

The utilization of the complete set of IGNS tools that were developed in relation to this project is already part of the routine procedure for thalamotomies and pallidotomies at the MNI. At the time this thesis was written, eight patients had undertaken the operation while the neurosurgeon consulted the supplementary information available from the atlas and its related tools.

Introduction

The two first Chapters are intended to provide historical and technical context to the work that was done in relation to that project. While Chapter 1 discusses the general field of Image-Guided NeuroSurgery, Chapter 2 focuses on the specific thalamotomy and pallidotomy operative procedures.

Chapter 3 presents the work involved in the creation of a deformable volumetric atlas of the basal ganglia for use in thalamotomy and pallidotomy specific IGNS. It also discusses the multiple deformation steps necessary to the implementation of atlas-based guidance in the Operating Room.

Chapter 4 and 5 constitute the main part of this thesis. The VIPER platform and its thalamotomy related features are presented in Chapter 4. Complementary work on the lesion modeling issue is the subject of Chapter 5.

Chapter 6 is divided in two parts. The first one describes in details the step-by-step procedure involved in the clinical utilization of the computer-guidance tools. The second part presents the clinical results obtained so far.

Chapter 7 briefly concludes the thesis and suggests ways of improving different aspects of the global project.

Chapter 1

Stereotaxy and Image-Guided NeuroSurgery

1.1 Conventional Stereotaxy

When a neurosurgeon performs an operation on the brain, he usually does so with one of two approaches: via an open craniotomy, or via a stereotactic procedure. Craniotomy is usually indicated for cortical neurosurgery. It consists of a partial and temporary removal of part of the cranial structure that allows the neurosurgeon to see the patient's cortex and to operate on it. In the case of deep brain neurosurgery, the structures of interest can only be reached with probes or high energy ionizing radiation beam. The neurosurgeon cannot see these structures during the operation. Therefore, some localizing device must be used to provide spatial information to the neurosurgeon. This is the goal of stereotaxy.

Stereotaxy is required for the following tasks:

- depth electrode placement within the brain for stimulation or for deep brain EEG recordings.
- approaching deep seated brain lesions with a high energy ionizing beam.

- performing thalamotomies or pallidotomies.

The localizing device (the stereotactic frame) must be rigid with respect to the patient's head so that the localization is reproducible and constant throughout the operation. Today's stereotactic frames are similar to those used in the first stereotactic procedures [27].

The first stereotactic procedure reported was performed on small animals by Horsley and Clarke in 1908, for depth electrode placement [22]. In 1947, Spiegel and Wycis [37] made the first use of a technique that could be applied to human stereotaxy, namely pneumo-encephalography which consists of taking images of the brain with x-rays after injection of air into the ventricles, normally filled with CSF. Since CSF has an absorption coefficient for x-rays that is similar to the brain, the ventricles can not normally be seen on a simple x-ray of the head. By injecting air in the ventricles, enough contrast is introduced to identify their positions on the x-ray film. Unfortunately this technique had the important clinical drawback of provoking strong headaches to the patient. Nevertheless this method was widely used and forms the basis for the technique that is used today, ventriculography, where a contrast-agent rather than air is injected into the ventricles (Chapter Two).

The representation of the brain that can be seen on x-rays is incomplete and ambiguous. This is easily explained by the fact that it presents a two-dimensional picture of an intrinsically three-dimensional physical volume. A point on this 2D image can be localized anywhere on the line (in physical space) projected onto that point (in image space). The advent of tomographic imaging techniques such as CT [23] and MRI [26], which provided complete three-dimensional representations of the brain eliminated this projection ambiguity. Assuming that the position of structures seen on images produced by CT or MRI can be related to the stereotactic frame coordinate system, complete correspondence between images and physical space may be obtained. The next section further discusses this issue.

1.2 Stereotactic Frame

The stereotactic frame used at the MNI consists of a cubic structure of bars mounted on top of a ring-like base. The different parts of the frame are made either of aluminum or rigid plastic, specifically assembled so that the aluminum parts do not form a closed loop; this is in order to avoid electrical current being induced in the metallic loop when the frame is imaged in a MRI scanner, which in turn induces artifacts. Plexiglas plates supporting fiducial markers can be mounted on the four sides and on top of the frame. For MRI, the fiducial markers consist of thin N-shaped tubes containing a solution of copper sulfate, giving a strong intensity signal when scanned. For CT, aluminum rods play the role of fiducial markers.

The frame is rigidly attached to the patient's head under local anesthesia using carbon fiber pins. Shallow holes drilled into the patient's skull allow the carbon fiber pins to be supported. Three to five pins are used to ensure a rigid fixation of the frame. Although there is a unique possible position for the frame to be re-mounted after removal, collars placed on the pins give added insurance that the re-fitting operation is reproducible. The scan is performed with the Plexiglas plates mounted on the plate so that the fiducial markers can be visualized and used for the registration between the images and the physical world (Fig. 1.1).

The idea behind the N-shaped configuration of the fiducial markers is trigonometrical. When performing a volumetric MRI scan with the frame mounted on the patient's head, the fiducial markers from each of the four surrounding Plexiglas plates (left, right, anterior and posterior) appear as a set of three bright dots on axial slices. The distances from the central dot (diagonal part of the tube) to the two other dots (vertical parts) can be used to compute the exact position of the central dot on the frame. On each axial slice, four of these points can then be identified and assigned a one-to-one correspondence between the physical frame space and the imaging MRI space (see Section 3.3).

Theoretically, if the imaging process was perfect, only three points would be needed to completely define the rigid transformation (3 rotations and 3 translations) from the imaging space to physical space. However, MR images are not perfectly accurate. Geometrical distortions due to

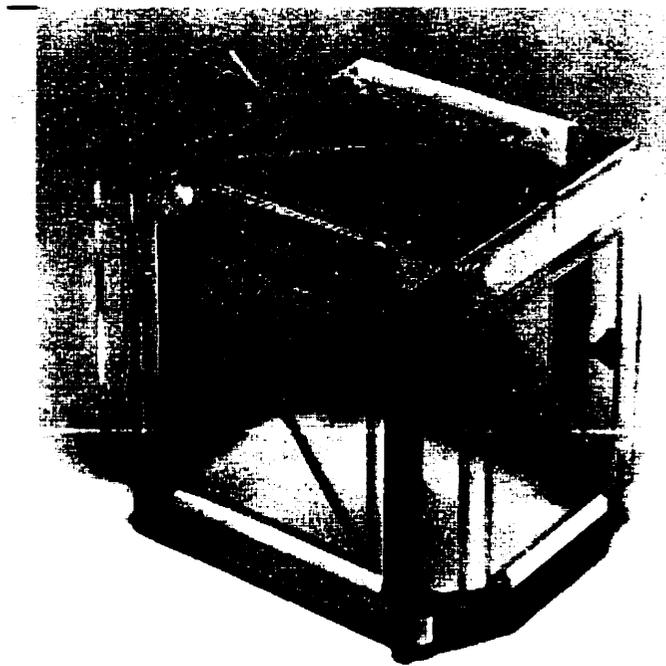


Figure 1.1: OBT (Olivier-Bertrand-Tipal) stereotactic frame with the fiducial marker plates.

field inhomogeneities and to other factors force us to use more than three points to compute the transformation. Instead, the point-to-point registration is usually done in four different planes, for a total of sixteen points. A least-square approach [1] is used in order to find the nine parameter transformation (three rotations, translations and scalings) that minimizes the error in the matching of the sixteen pairs [36].

An arc system can be mounted on the frame itself to support needle-like neurosurgical tools (Fig. 1.2). The principle of the arc system resides in its angular-independent aiming characteristic. The arc itself can be rotated around its axis, while the tool holder on the arc can be fastened at any position on the arc. The possible positions of the tool holder are thus described by part of a sphere. By properly setting the position of the arc with respect to the frame so that the centre of the sphere is exactly at the point of interest (the target), the neurosurgical tool will always point towards this target wherever it is attached on the arc system.

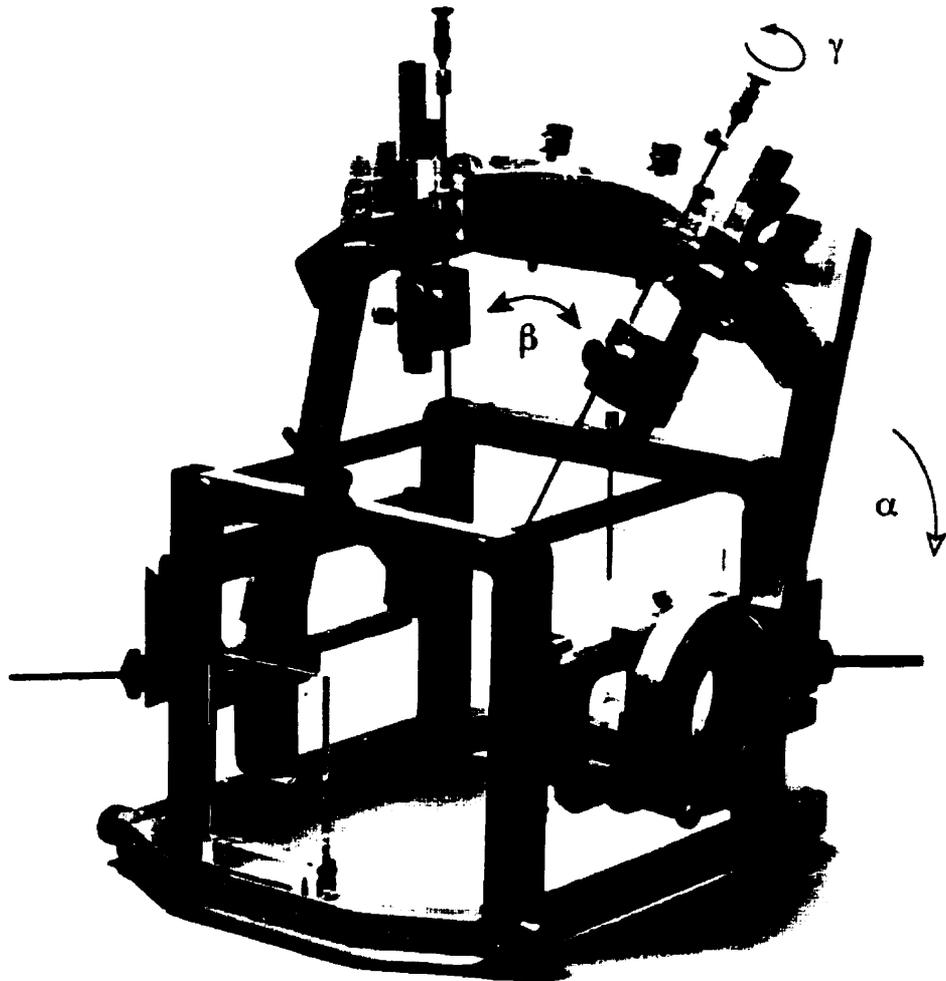


Figure 1.2: Arc system mounted on the stereotactic frame.

Two biopsy needles are attached to the arc system installed on the stereotactic frame. The arrows indicate the declination (α), azimuthal (β) and twist (γ) angles.

1.3 Image-Guided Neurosurgery

Stereotaxy was developed to provide the neurosurgeon with a localizing device for deep brain procedure where it was impossible to visualize the structures by eye. In principle, this localizing device should not be needed for cortical procedure where craniotomy is performed, since the neurosurgeon can directly see the brain structures. However, some difficulties arise even in these procedures:

- The region to be operated is not always found exactly at the surface of the cortex, so critical structures and vasculature might be in the way.
- The surgical site may have the same appearance as its surroundings, which must be spared.

The goal of *Image-Guided NeuroSurgery (IGNS)* is to provide guidance to the neurosurgeon by presenting images of the brain that are properly registered with the physical space. This means that a complete correspondence between features appearing on the image and position of those features in the actual physical space (the patient's head) must be provided. These images can come from a number of different modalities, such as MRI or CT scans, vascular images (DSA, MRA) [20, 19], functional scans (PET, fMRI) [30, 32], chemical concentration scans (MRSI) and others (atlases, EEG...).

For deep-seated regions of the brain, the stereotactic frame provides the correspondence between image and physical spaces. For open-cranium procedures, it is awkward if not impossible to attach a frame to the patient's skull, and so the correspondence between the two spaces must be defined by another localizing device. At the MNI, an IGNS system is used to guide the surgeon. It consists of an articulated arm (FARO Industries, Florida) (Fig. 1.3) linked to a computer to transfer 3D positional information from the tip of a probe into the computer. Other localizing devices based on mechanical [17, 16, 42], optical, ultrasonic [2] or magnetic [24] principles have been developed as well. The articulated arm used at the MNI is a mechanical localizer, made of aluminum shafts joined together so that they can rotate with respect to one another. A total of six degrees of freedom are available, allowing relatively free movement of the 20cm probe that is attached to the end of the arm. The FARO arm is a passive device in the sense that it is not



Figure 1.3: The FARO Surgical Arm.

motor-controlled but rather manipulated by the neurosurgeon to the desired position. The arm sends information on the relative position of the different parts (shafts) of the arm through angle potentiometers found in the joints. These potentiometers send analog signals to an analog-to-digital converter interfaced to a computer. The computer calculates the three-dimensional coordinates of the tip of the probe and the orientation vector of the probe itself from this information, and then sends it to the navigation console.

A typical IGNS system consists of an image display that presents the probe position, related to MR and CT images. At the MNI, the ISG Viewing Wand (ISG Technologies, Mississauga, Ontario) is used on a routine basis as an IGNS system for craniotomy operations (Fig. 1.4). The images can be visualized as 2D slices through the volume under different orientations (coronal, sagittal, axial and along the trajectory of the probe), and as 3D objects that have been pre-segmented (pre-identified) on the volumetric data set. A detailed discussion about 3D visualization will come in Chapter 4.

Since these operations are performed without the stereotactic frame, no fiducial markers appear on the scans to allow for the registration of image space and physical space. Registration

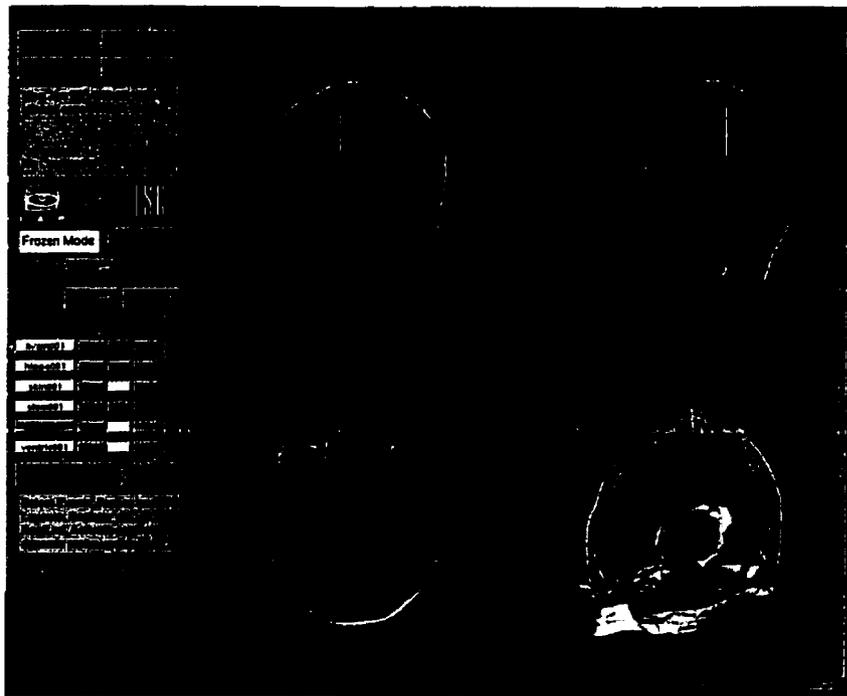


Figure 1.4: The Viewing Wand IGNS system main screen.

is achieved in these cases through the FARO Arm, which links the computer to the real world. The registration is performed through the following steps: while the neurosurgeon holds the articulated arm in a position where the tip of the probe touches an external landmark on the patient (e.g. the bridge of the nose), the computer assistant uses the computer mouse to move the screen cursor at the corresponding point on the image. When both the tip of the probe and the screen cursor are placed at corresponding positions, registration of that landmark is made. This is performed for four or five landmarks after which good agreement between the images and the physical space is usually obtained. More sophisticated methods (surface matching) can be used if this kind of registration does not give satisfactory results.

Once a proper registration is achieved, the neurosurgeon can navigate through the images by moving the tip of the probe to different positions. In the 2D views, the software shows slices through the volume at the tip of the probe, while the 3D-view shows a rendered version of the probe relative to pre-segmented 3D objects.

Since some additional developments to the platform were done by our research group, the Viewing Wand version available at the MNI contains features that are not standard on the commercial version. These features includes integration of PET, MRA and DSA imaging modalities and stereoscopic display of 3D objects [20, 19, 34, 33]. Additional improvements to the capabilities of the software have since been implemented. A more recent IGNS system, developed entirely in the MNI laboratory, is discussed in Chapter 4.

Chapter 2

Thalamotomy and Pallidotomy

This chapter discusses surgical procedures used to alleviate the symptoms of Parkinson's syndrome.

2.1 Symptomatic Parkinson's syndrome

Symptomatic Parkinson's syndrome manifests itself in a variety of ways, including tremors, rigidity, bradykinesia, impaired gait, dysphasia and ocular disturbance. Most of these symptoms can be relieved by the prescription of L-dopa, except for tremors, which are more resistant to the drug. In some patients, the tremors are even made worse with L-dopa. Tremors result from "rhythmic alternating contraction of opposing muscle groups usually of the distal extremities" [40]. In the case of Parkinson's syndrome, the tremors are known to come from disorders in the extrapyramidal motor system. These tremors are considered to be slow (about 4-5 oscillations per second) in comparison with other types of tremors that can go up to 20 oscillations/s. The frequency of oscillation for parkinsonian tremors is not constant: it can vary from as slow as 2 osc./s up to 7 osc./s.

The most common manifestation of these tremors is found in the hand muscles, where "the alternating contractions of flexors and extensors of the fingers, of abductors and adductors of

the thumb produce a pill-rolling or bread-crumbling movement, which is often associated with a rolling movement of the forearm due to the alternating contractions of pronators and supinators” [40]. The tremors disappear at sleep, and come back a few minutes after awakening. They are mostly present when the arms (or legs) are at rest. They increase in amplitude when the patient undergoes emotional or physical stress.

Parkinsonian patients suffering from significant tremors and those who have gained only a mild improvement (or no improvement at all) of their general state from L-dopa are good candidates for thalamotomy and pallidotomy. At the MNI, pallidotomies are usually performed when the patient suffers from rigidity more than from the tremors. Ideally, candidates for thalamotomy and pallidotomy should be under 60 years of age, with unilateral tremors or rigidity. The probability of almost complete relief of symptoms (without side-effects) for these candidates is around 80% [41].

2.2 History and description of thalamotomy and pallidotomy

In chapter 1, we mentioned the work of Spiegel and Wycis (in 1947) on the development of the first stereotactic procedure on man [37]. Around the same period, Meyers was studying the effect of excision of the head of the caudate nucleus and of the severance of the ansa lenticularis [29]. He demonstrated that both of these procedures could greatly reduce the tremors and lessen rigidity of parkinsonian patients, inviting neurosurgeons to focus their attention on the basal ganglia for cure of motor dysfunction.

Hassler was the first to suggest the nucleus ventralis lateralis as a target for the relief of dyskinesia, based on the extensive connections between this part of the thalamus and the motor cortices [18]. Currently at the MNI, the specific target for thalamotomy is the *nucleus ventralis intermedius (Vim)*, while the target for pallidotomy is part of the *globus pallidus medialis (Gpm)* and sometimes a small part of the *globus pallidus lateralis (Gpl)* (Fig. 2.1). However, it is difficult to identify exactly the position of these targets with conventional imaging tools, such as ventriculography (Section 2.3). Even though the midplane and anteroposterior commissural (AC-



Figure 2.1: The Vim node of the thalamus and the Gpm.

The Vim node is shown in dark grey, and the Gpm in light grey (from the Schaltenbrand and Wahren atlas [35]).

PC) plane can easily be localised on ventriculograms, the distances at which the target structure lies from these planes change from one patient to the other, mostly because of the variable size of the third ventricle. The distance between the third ventricle and the target also depends on the size of the former.

In order to address this problem, physiological verification of the position of the target with respect to the stereotactic frame coordinate system is performed prior to the excision itself. The position of the target is confirmed through electrical stimulation of the basal ganglia. An electrode is inserted into the patient's brain through the same burr-hole used for the injection of the ventriculography contrast agent. The neurosurgical tool itself is called a stimulator and consists of a long metallic shaft through which runs a small conducting wire. The wire is attached at one extremity to a screw that controls the extent of the protrusion of the wire from the tip of the shaft. The stimulator currently used at the MNI extends a maximum of 14 mm, obliquely to the shaft (Fig. 2.2). The wire is connected to an oscilloscope that controls the electrical stimulus. The waveform and frequency of the electrical stimulus are constant, while the voltage can vary from

0.25V to 2V.

While the neurosurgeon performs the stimulation, observation of sensory, motor or visual responses from the patient are noted [5]. Motor responses come from stimulation of the internal capsule, an important white matter structure that should be spared from the excision. The internal capsule defines the border of the Vim on the lateral side, and of the Gpm on the medial side. Sensory responses come from Vci and Vce nuclei in the thalamus, and these nuclei should be spared as well. They both touch the posterior part of the Vim node. From the information obtained with the electrical stimulation, the neurosurgeon is able to determine the position of the target relative to the frame coordinate system with acceptable precision.

Once the surgeon has obtained enough information about the position of the important deep brain structures of the patient, he can proceed with the curing lesion through excision of the Vim or the Gpm. The lesion can be created in a number of ways: electrolytically, mechanically, by injection of corrosive fluids, thermogenically (using radio-frequency fields), cryogenically, ultrasonically or by radioactive isotope seeds. Predictability of the size of the lesion is a concern with electrolytic and injection lesions. Thermogenic and cryogenic lesion can lead to haemorrhage if great care is not taken. Ultrasonic equipment is quite expensive. Radiation (tele- or brachy-therapy) and mechanical devices are probably the two best choices for this task, though other methods are employed as well. At the MNI, the lesions are created using a mechanical device (leukotome) consisting of a narrow metallic wire protruding from a shaft (similar to the stimulator) so that it forms a loop which acts like a knife when the tool is rotated around its shaft (Fig. 2.2). Once again, a screw at the other extremity of the shaft controls the extent of the loop. The leukotome extends to a maximum distance of 7mm out of the shaft.

The main advantage of the leukotome over other tools resides in its ability to perform lesions of various shapes. While other methods are usually limited to spherical lesions, the leukotome allows for pre-defined lesions that can be geometrically adapted to the target volume to be excised. The rotation of the leukotome around its shaft does not need to be a full 360 degree with the wire out of the shaft by a constant extent, but can rather be divided in many partial rotations with different protrusions of the wire. Provided that the neurosurgeon knows exactly the

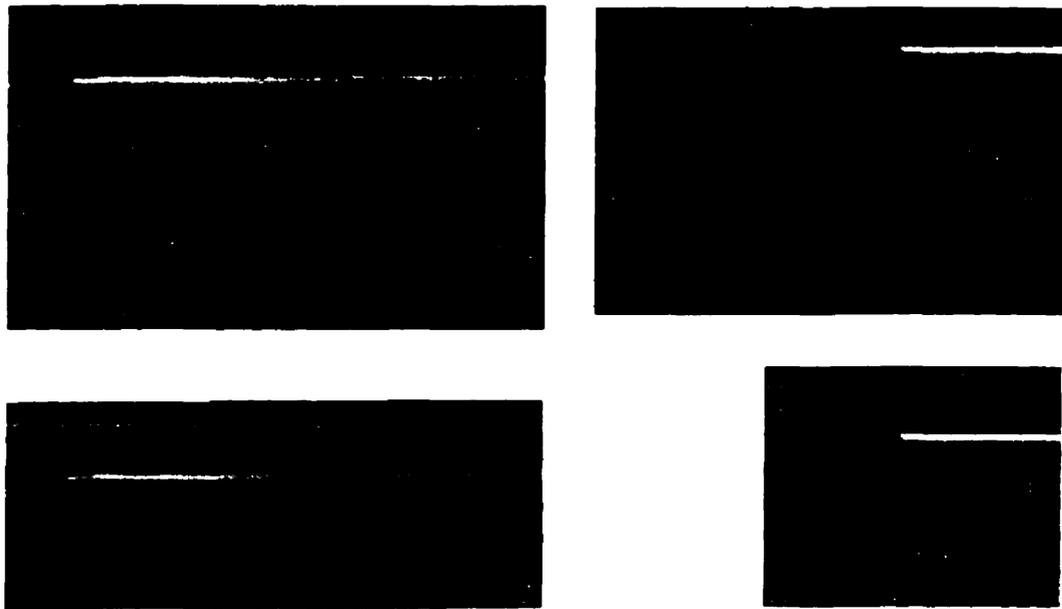


Figure 2.2: Optical scans of the leukotome and the stimulator with different protrusions of their wires

geometry of the region that needs to be excised, a precise tailoring can be achieved by creating many small lesions of appropriate size in different angular regions (sectors) around the shaft of the leukotome. Either eight or sixteen sectors of 45 or 22.5 degrees respectively are used to define the tailored excision. Other means of creating lesions (i.e. hyperthermic or cryogenic approaches) offer the advantage, in theory, of reducing risks of bleedings. On the other hand, the lesion margins are not as easy to identify.

2.3 Conventional imaging modalities for thalamotomy and pallidotomy

The standard procedure at the MNI for localization of the target point where the tip of the leukotome should be placed includes a conventional stereotactic imaging technique called ventriculography. The principle of this technique is to enhance the contrast (on x-ray films) between ventricles and surrounding brain tissue by injecting a contrast agent into the intra-ventricular

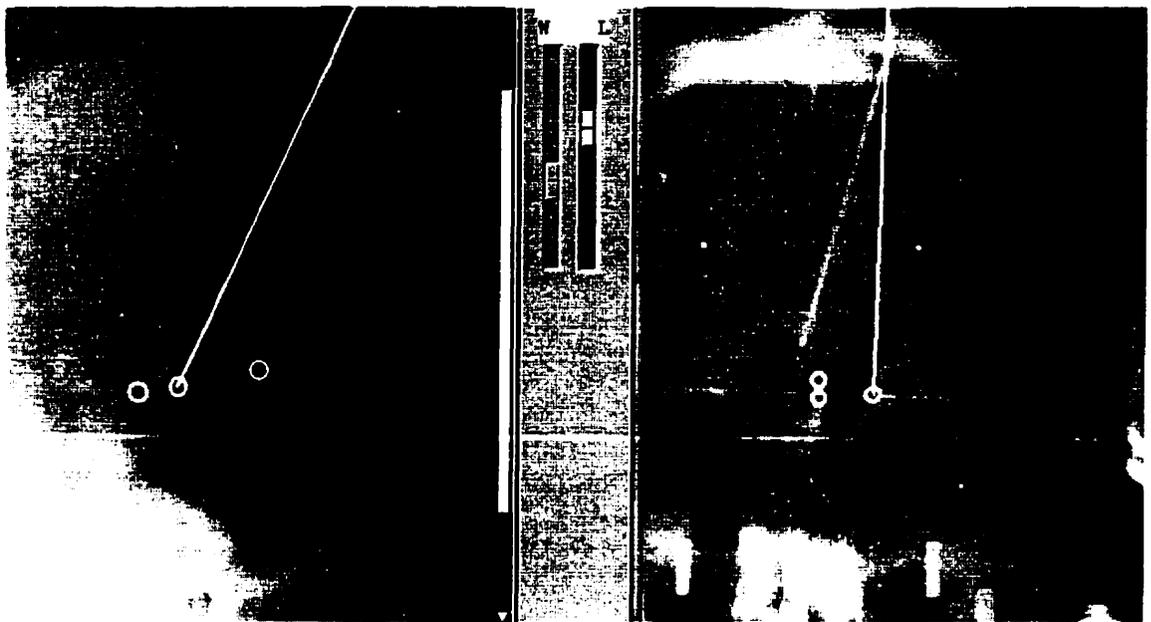


Figure 2.3: Ventriculograms.

An example of lateral (left) and antero-posterior (right) ventriculograms with the AC, PC and target points identified.

cerebro-spinal fluid (CSF). Since CSF, white matter and grey matter have similar attenuation coefficients for x-rays, very little contrast appears on x-ray films of the head. The contrast agent brings out the ventricles on films, and allows the neurosurgeon to identify the AC and PC points.

Two x-ray films (called ventriculograms, one antero-posterior and one lateral) of the head are taken with the contrast agent in the ventricles and the frame attached to the patient's skull (Fig. 2.3). Fiducial markers on the eight corners of the cubic-shaped frame appear on both ventriculograms. From the projection of those markers onto the films, the three-dimensional position of a point appearing on both films can be calculated. A standard manual method using the triangulation principle of the projection is still in use today and constitutes the gold standard for localization of the AC and PC points in the 3D frame coordinate system. A program has been developed at our lab [12] that performs an automatic calculation of the two homogeneous perspective projection matrices from the two-dimensional position of the fiducial markers on the film. The program then allows the user to obtain the three-dimensional position of a point which is identified on both films, so that identification of the AC and PC points can be achieved more

easily. Use of this program is now part of the standard procedure, but a manual calculation is still performed to verify the computer's output.

During the operation, the stimulator and leukotome position in the patient's head can only be verified by intra-operative x-rays, but the contrast agent that was injected for the ventriculography has usually washed out by this time. In order to compare the position of the neurosurgical tools while they are inside the patient's head with respect to the ventricles, an intra-operative x-ray film can be superimposed on the ventriculogram. However, the projection geometries of the intra-operative x-ray and the ventriculogram must be exactly the same. This implies that the exact position of the x-ray tube with respect to the patient and the film when the lateral ventriculogram is taken must be accurately reproducible. At the MNI, this is done by attaching the x-ray tube and the cassette holder to the OR bed.

Chapter 3

Atlases and Automatic Labeling of the Deep Brain Structures

In chapter 2, we presented the standard procedure for thalamotomy and pallidotomy at the MNI. We will now discuss the weaknesses of that procedure and propose some solutions.

The goal of thalamotomy and pallidotomy is to cut out the fibers linked to the Vim node of the thalamus, or the ones linked to the Gpm. The relation of these two target structures to their neighborhood is complex in its geometry. The two structures are encapsulated in an environment of other structures, some of which are fundamental to important functionality of the deep brain. The internal capsule, for instance, is one such critical structure that should absolutely be avoided during the excision. It is a white matter structure that conducts information between the deep brain and the cortex. Its geometry is very complex, and its extent covers a large portion of the deep brain. The only references that the neurosurgeon can look to during the operation (in order to localize the position of the internal capsule) are books containing atlases of the anatomy of the deep brain [38]. These books are usually presented as a set of slices through the brain with outlines of the different structures in place. These atlases are created from studies performed on brains that are different from the patient's brain at least in size and sometimes in geometry. The orientation of the patient's brain with respect to the slices presented in the atlas as well as the path followed by the neurosurgical tools with respect to the structures described in the atlas are

difficult to understand. Hence we need to address the following problems:

- Anatomical inconsistencies between atlas and patient.
- Undefined orientation of the patient with respect to the atlas.
- Unknown position of the neurosurgical tools with respect to important structures described in the atlas.

A quick solution to all of these problems would be to find an imaging modality that allows the neurosurgeon to easily discriminate between all the important parts of the deep brain anatomy. These structures consist of either white or gray matter, or sometimes both. CT scans show very little contrast between white and gray matter in comparison to MRI scans where white and gray matter can be differentiated easily. The problem is that two regions of the anatomy responsible for different functions of the brain can appear the same on an MRI scan since they have very similar white and gray matter content. This is the case for the Vim, Vci and Vce nodes of the thalamus. Still, MRI constitutes the imaging modality that gives the best information on the anatomy of the thalamus and its neighboring structures. Moreover, MRI data sets are presented in a volumetric format; this means that a complete 3D volume (the patient's brain) is represented as intensity values stored in individual cubes (voxels). This format can be visualized in a number of ways that will be explained later.

The solution we propose is an enhancement of the information already available from routine MRI scans. We want to incorporate the precise outlines of an anatomical atlas of the deep brain to an MRI scan performed on any patient. In order to solve the problems we stated earlier, we need this atlas to be volumetric, deformable and easily orientable with respect to the patient's brain. The positioning of the neurosurgical tools is discussed in section 4.3.

3.1 Schaltenbrand and Bailey Atlas

At the MNI, the Schaltenbrand and Bailey atlas is the principal anatomical reference for neurosurgeons performing thalamotomies and pallidotomies. This book gathers together many differ-

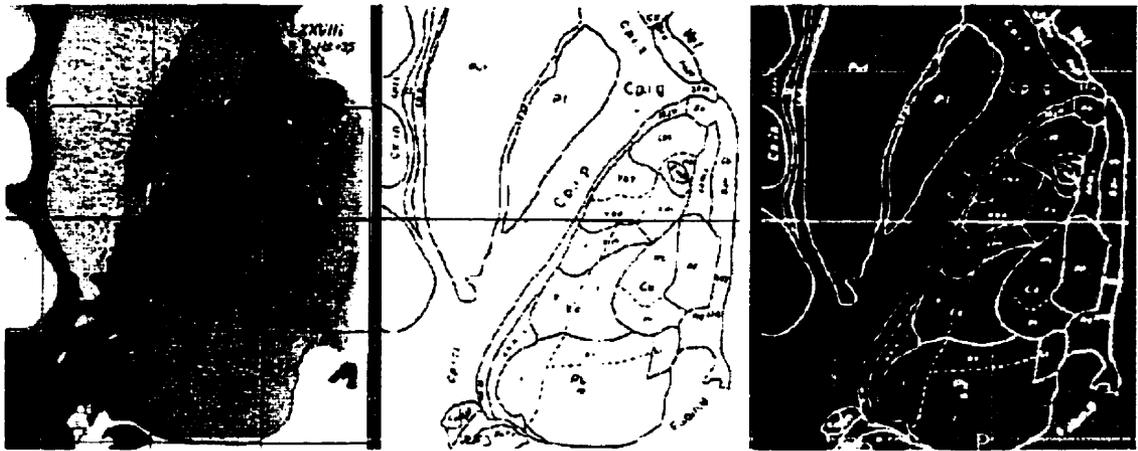


Figure 3.1: Schaltenbrand atlas.

One axial slice from the Schaltenbrand transverse series: anatomical stained slice (left), transparent overlay contours (middle) and superimposition of the two (right).

ent data sets. For our purpose, three of those sets are of interest: the coronal (plates 35-40), sagittal (plates 41-51) and transverse (plates 52-57) cryogenic slices of cadaver's brain [35]. Sagittal and axial sets both come from the same brain (LXXVIII) (from the left and right hemispheres respectively). The coronal set comes from another brain (LXVIII). Each of these three data sets consists of a series of micro-thin cryogenic slices that were stained to discriminate between structures presenting different cyto- architectures. The slices are not necessarily contiguous in space; each slice is separated from the next one by a certain distance varying between 0.5mm up to 3mm. Photographs of these slices were then shot under constant geometry. The outlines of the anatomical regions were drawn on acetate sheets overlaid on top of the photographs. The atlas series is presented on large paper sheets with four photographs per sheet. An acetate sheet on which the outlined regions are drawn is associated with each photograph sheet (Fig. 3.1).

Our goal is to incorporate the information found in the atlas into an MRI data set. The merging of the atlas with the MRI should be done on a computer workstation rather than on printouts to allow for better interactivity with the two data sets. The difficulty is to create an easily deformable volumetric atlas from an atlas made of two-dimensional slices with varying inter-slice separations. Since the three series of interest to us are in principle independent from one another, we had to decide which one should be used as the initial data set from which the volumetric atlas

would be created. It is clear that the volumetric version presents more consistent images when visualized as slices taken in the orientation of the original data set (the Schaltenbrand atlas set) than in the two other directions since the interpolation in the third dimension would reveal inherent misalignment of the original data set. We decided to use the transverse series, since the extent of the structures in this plane was larger than in the two others, and also because the neurosurgeon was more familiar with them. In this series, about a hundred structures are outlined. We limited the construction of the volumetric atlas to the sixteen regions that were of interest to the neurosurgeon in the thalamotomy and pallidotomy procedures.

The steps involved in the creation of the volumetric version of the Schaltenbrand atlas are the following:

- Scanning of both the cryogenic slices and the transparent overlays into the computer.
- Stacking (concatenation) of the individual slices into a 3D data set.
- Extraction of the two-dimensional contours (from the transparent overlay) in a vector format.
- 3D interpolation of the two-dimensional contours.

The scanning of the cryogenic slices and of the acetate overlays was done on a HP ScanJet IIcx scanner connected to a Macintosh computer (Apple Computers, Cupertino, California). In principle, only the acetate overlays need to be scanned since we want to use the contours to create the volumetric version of the Schaltenbrand atlas. But the cryogenic slices present a grid drawn on the photographs for localization and alignment of the different slices. Hence these images were also scanned in order to perform proper stacking of the acetate contours afterward.

Once the scanning of the slices was performed, we aligned the acetate contours with respect to their corresponding cryogenic slice using Adobe Photoshop (Adobe System Inc., San Jose, California). This alignment is necessary since the stacking is performed using the cryogenic slices instead of the acetate. The alignment is performed by visual inspection of the contours with respect to the underlying anatomy shown on the cryogenic slices. The acetate contours are

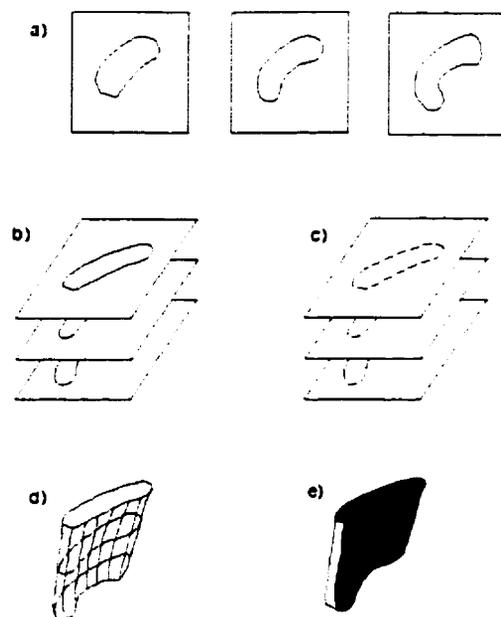


Figure 3.2: Construction of the volumetric atlas.

a) Scanning, b) Stacking, c) Vectorizing, d) Interpolating and e) Final volumetric version of the structure.

represented in Photoshop as black lines on a transparent background, and can be moved over the cryogenic images as an overlay. The acetate contour is then saved under the same geometry as its underlying cryogenic slice.

We created a program (written in C) to perform the stacking of the slices in a volumetric data set. At the MNI, the MINC^a format is used as a uniform environment for manipulation of volumetric data sets such as MRI and CT scans. We stored the stacked version of the 2D slices in this format. The program requires 3-D position of each of the individual slices to be accurately specified to perform the stacking. For each slice, the distance from the transverse plane to the AC-PC plane must be given, as well as the pixel coordinate of one point in every image lying on an axial line through the registered image stack. We arbitrarily chose that point at the crossing of two grid lines. The resulting volumetric data set consists of zero intensity voxels where no transverse data is present in the atlas set. When visualized in tri-planar display, this volumetric data set presents

^a The MINC format is an adaptation of the general data format NetCDF (from the UCAR - University Corporation for Atmospheric Research). MINC stands for Medical Image NetCDF.

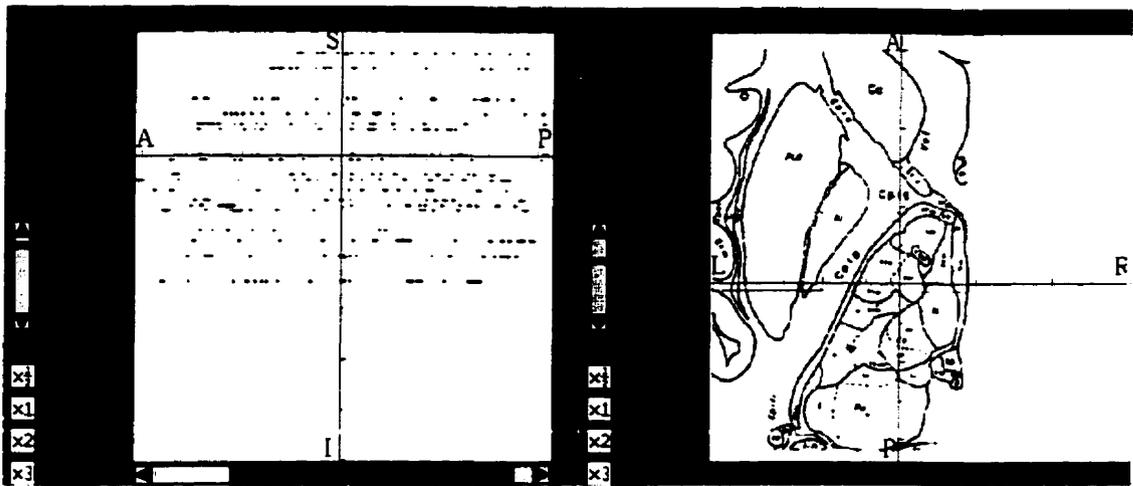


Figure 3.3: Two views of the stacked slices.

The sagittal view (left) presenting a rough set of dots, showing the need for a volumetric interpolation between the Schaltenbrand axial slices.

either acetate contours or a completely black image in the transverse view, while a set of sparse dots can be seen in the two other views (Fig. 3.3). In order to convert these data into a full 3-D data set, we must perform an interpolation operation between the two-dimensional slices. This interpolation operation was performed using spline parametrization of 3-D curves [15]. For our purpose, we chose Hermite cubic interpolation from the large number of spline parametrizations found in the literature. Hermite cubic parametrization is well designed to describe a relatively smooth curve through a set of ordered points in space, and is probably the simplest way of passing a smooth curve through points. The parametrization of a curve going from one point to the next in the list is mathematically defined by the following set of three equations:

$$\begin{aligned}
 x(t) &= a_x t^3 + b_x t^2 + c_x t + d_x \\
 y(t) &= a_y t^3 + b_y t^2 + c_y t + d_y \\
 z(t) &= a_z t^3 + b_z t^2 + c_z t + d_z
 \end{aligned}
 \tag{3.1}$$

where x , y and z are the standard coordinates in three dimensions, t is a parameter that runs

from zero to one and $a_x, b_x...$ are the coefficients that describe the behavior of the curve between two points. These coefficients are determined by imposing certain constraints on the curve. Since four parameters must be defined for each of the three equations, we need four constraints in each dimension. The two first constraints are the most obvious ones: the curve must start at one point in the list and end at the next point. This means that at $t = 0$, $x(0)$, $y(0)$ and $z(0)$ must be equal to x_i , y_i and z_i when the curve runs from point i to point $i + 1$. Then at $t = 1$, $x(1)$, $y(1)$ and $z(1)$ must be equal x_{i+1} , y_{i+1} and z_{i+1} . The two remaining constraints are related to the *speed vector*^b the beginning and end points. We ask for the speed vector at point i to be parallel to the vector joining points $i - 1$ and $i + 1$. This results in smooth-shaped curves passing through our set of points. The magnitude of the speed vector sets the smoothness of the curve. It was rather arbitrarily defined as one third of the magnitude of the vector running from point $i - 1$ to point $i + 1$. This factor resulted in satisfying interpolation.

Cubic parametrization in each dimension allows for the curve to move in any direction. For our needs, it is useless for the parametrization in z to allow for non-linear interpolation between two points. For instance, we do not want the curve to describe loops in the z direction. In fact, we need this parametrization to be linear so that we can retrieve the x and y coordinate of the interpolation between two points at any given z to create the missing slices. By forcing the z parametrization to be linear ($z(t) = t + d_z$), any t value corresponds to the fraction of separation in z between the two interpolated slices.

Practically, each curve must be described as a set of points through which the curve passes. This is called a vector format. A C-program (written by Diego Clonda) retrieves the vector format given a two-dimensional image and the coordinate of one point inside the curve to be vectorized. Once this is done for all sixteen structures in every plane in which they are outlined, the vector formats are fed into another program that perform the Hermite interpolation (in vector format) between original slices, at intervals of 0.25mm. This complete set of vector formats (defined at every 0.25mm) is finally sent to a last program that creates a volumetric version of each of the sixteen individual regions (in the MINC format). At this point, we have a complete

^b The speed vector is defined as the first derivative of the parametrized curve at a given point on the curve, and has the characteristic of being parallel to the curve at that point.

volumetric version of the sixteen regions of the Schaltenbrand atlas in which we are interested.

3.2 Model MRI Data Set: Superbrain

The benefit of having a computerized volumetric version of some structures defined in the Schaltenbrand atlas is still somewhat limited from a neurosurgical point of view. Anatomical variability from one patient to another prevents a simple overlap of the atlas on the MRI in proper orientation (i.e. after applying the proper rotation and translation to the atlas) to be accurate enough for IGNS needs. This overlap is not precise enough to be reliable. Even after proper scaling in each individual direction, local anatomical variability is not taken into account even though it might be sufficiently great such that the atlas regions do not correspond to the actual anatomy they represent. Hence we need not only to rotate, translate and scale the atlas, but also to deform it locally (in a non-linear fashion) so that it can be adequately superimposed onto the patient's MRI with the different structures appearing in the correct place.

Our objective is to have a non-linear deformation that may be performed on a routine basis without any intervention of specialized staff (a neurosurgeon or a neuroanatomist). Dr Louis Collins has developed an algorithm that performs this kind of non-linear deformation between two MRI data sets. The *ANIMAL*^c algorithm is a 3-D cross-correlation technique that uses a multi-scale approach [10, 11, 9]. It first computes the best global linear transformation, and then defines the local deformations on finer and finer scaling. This kind of non-linear deformation corresponds to our need, except that the algorithm requires MRI images as both the original and target volumes.

In order to use this algorithm to warp the atlas onto the patient's MRI, we must first go through an intermediate step. Thin-plate spline 3D interpolation [7] allows us to perform a smooth non-linear deformation from any volumetric data set to another provided a set of point pairs. Each pair of points consists of corresponding anatomical landmarks, one point being found in the MRI data set and the other on the atlas. The strategy is to manually identify corresponding landmark

^cAutomated Nonlinear Image Matching and Anatomical Labelling.

on the atlas and on a model MRI data set one time only, so that the atlas could be definitively warped onto this model MRI. Afterward, the non-linear transformation from the model MRI to any patient's MRI could be automatically computed using the ANIMAL algorithm. Since the atlas is already in correspondence with the MRI model space (from the manual registration), this last transformation can be applied to the atlas in order to bring it onto the patient's MRI. Let us summarize the steps involved in bringing the atlas into patient's space:

- Manual tagging of corresponding landmarks on the atlas and on a model MRI (performed once and for all by a neuroanatomist)
- Deformation of the atlas onto the model MRI using thin-plate spline interpolation from the set of corresponding landmarks (once and for all).
- Automatic calculation of the transformation from model MRI space to patient's space (for every patient).
- Deformation of the model-space atlas onto the patient's MRI space (for every patient).

Only the two last steps (which require no intervention from a specialist) must be performed for every individual case.

The choice of the model MRI is a fundamental element in the success of all these steps. The quality of the manual tagging of the atlas and the performance of the ANIMAL algorithm depends strongly on this choice. The most important criterion is the signal to noise ratio (SNR) in the image. Lower noise in the model data set will result in a better non-linear fit and easier manual tagging. Dr Colin Holmes [21] created a volumetric MRI data set presenting a high SNR. It constitutes an average of 27 scans of his own head. The noise in the data set is inversely proportional to the square root of the number of acquisition, so that this data set (that we call the Superbrain from its low-noise characteristic) reveals an SNR about 5.1 times higher than one obtained with a standard volumetric acquisition under this same protocol (Fig. 3.4).

The manual tagging of the atlas to the model MRI was performed by Dr. Noor Kabani, a neuroanatomist at the MNI. A total of 250 corresponding landmarks were identified to ensure

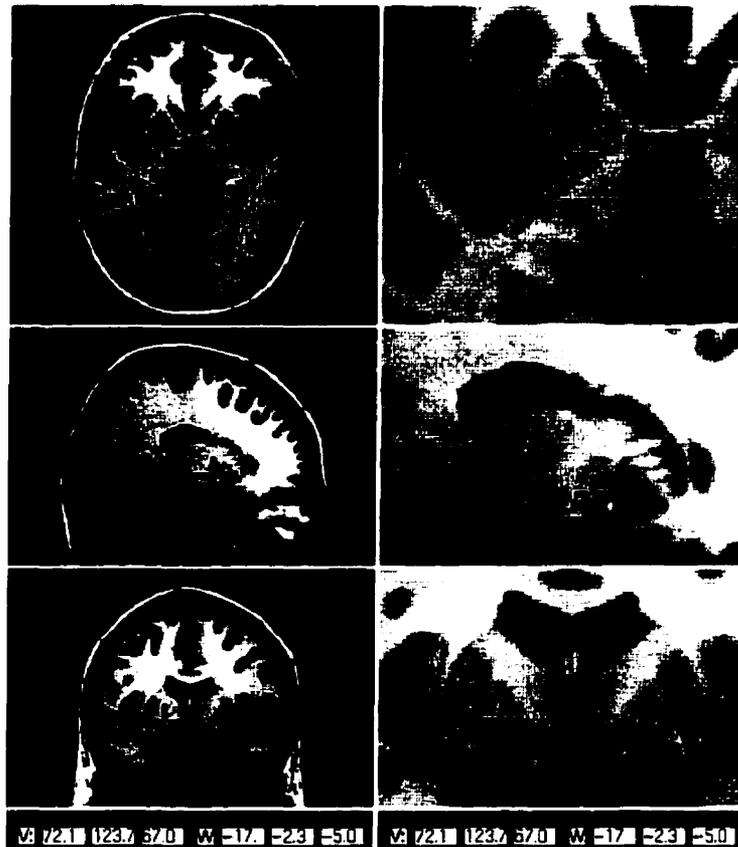


Figure 3.4: Our MRI model data set: Superbrain



Figure 3.5: Part of the tagging performed between the Atlas (left) and the MRI model (right).

that the atlas precisely overlapped with the model MRI (Fig. 3.5). The thin-plate spline interpolation forces all 250 points from the atlas to be mapped onto their corresponding points on the model MRI, while the regions in between the points get smoothly warped into the new space.

3.3 Patient space

The warping of the model MRI (and of the atlas) onto the patient's brain (using the ANIMAL algorithm) is an automatic procedure that requires a lot of CPU time (specifications). In principle, this should not be considered as a limitation since the algorithm can run overnight all by itself. However, it happens that the complete operative procedure takes place over one day only, so that the MRI scan with the frame in place is acquired in the morning whereas the excision itself must be performed during the afternoon. In those cases, there is not enough time to compute the non-linear transformation. Fortunately, the standard protocol for thalamotomy and pallidotomy includes a pre-operative MRI scan acquired the week before the operation. We use this data set to compute the non-linear transformation. Then, on the day of the operation, a simple linear reori-

entation (translation and rotation) must be calculated to go from pre-operative to intra-operative space, since both scans come from the same head. We assume that the brain does not move or deform inside the skull between the two scans. The linear reorientation can be performed automatically on the computer and requires little CPU time (15minutes).

Up to now, we have discussed how to bring the atlas onto the patient's MRI without specifying exactly how this could be used in the OR. In order to be considered as useful data sets, both the atlas and the patient's MRI must be navigable in a coordinate system understood by the neurosurgeon. Thus they must be registered with respect to the actual physical position of the neurosurgical tools. As stated in Chapter Two, these tools are rigidly attached to the stereotactic frame that was screwed into the patient's skull prior to the operation. Rulers drawn on this stereotactic frame define the coordinate system that the neurosurgeon refers to when localizing the target. Hence the atlas and the MRI of the patient's brain must be transformed into that frame space.

The transformation to frame space is obtained as explained in Chapter 1, using the Z-shaped fiducial markers that appear on the MRI scan. By identifying all four markers (left, right, anterior, posterior) in a plane for four different planes, the rigid transformation from the MRI space to the frame space can be evaluated to a reasonable degree of precision [1, 36]. Since the atlas has already been transformed to the MRI space, this last rigid transformation can be applied to it. Now the neurosurgeon can verify what structure will be stimulated or excised (according to the atlas) once the neurosurgical tools are inserted deep into the patient's brain. We explain in the next chapter how this virtual navigation through the brain can be performed in the operating room.

3.4 On the Different Spaces and Coordinate Systems

A number of different spaces have been presented up to now. Here we review their relationship to one another. Note that except for the frame space (which differs for historical reasons), all other coordinate systems are defined as X-axis running from left to right, Y-axis from posterior

to anterior and Z-axis from inferior to superior.

3.4.1 Atlas Space (or Schaltenbrand Space):

The coordinate system used by Schaltenbrand for the transverse series is relative to the brain that was used in the creation of that series. It is based on the AC-PC line and the mid-plane separating the two hemispheres. The origin of the system is localized at the anterior commissure (AC), and the scale is in millimeters.

3.4.2 Model MRI Space (or Superbrain Space):

The origin is found close to the anterior commissure (AC), and the coordinate are in millimeters.

3.4.3 Native MRI space (Pre-op, Intra-op, Post-op):

A coordinate system is associated to each one of the three scans that is standardly performed on the patient to be operated. In all three cases, the origin on X and Z-axis is located in the centre of the image, while it is found at the most posterior part of the image for the Y-axis.

3.4.4 Frame Space:

The OBT Frame space is completely different from all the others. It is a left- handed coordinate system with the X axis running from anterior to posterior, the Y axis from inferior to superior and the Z axis from left to right. The origin of the three axis is the midpoint of the anterior section of the base ring. Numbers found on the rulers drawn on the frame represent centimeters.

3.4.5 Physical Space (Brain Space, Patient Space):

All these terms are used alternatively to describe the “real” physical space in which the patient’s head lies (in the OR). Identification of the target (in frame coordinates) by the neurosurgeon can

also be seen as a registration between the brain space and the frame space. The position of a given structure or landmark in brain space is usually given with respect to the AC PC line, exactly as it is defined in the atlas space.

Chapter 4

Image-Guided NeuroSurgery Tools for Thalamotomy and Pallidotomy

A number of graphical interfaces for neurological navigation and for IGNS are already commercially available. At the MNI, the Viewing Wand from ISG Technologies is used on a routine basis for pre-operative and intra-operative imaging needs in neurosurgery. This platform provides standard tri-planar views of either MRI or CT scans, as well as trajectory views. It is also possible to visualize 3D surface-rendered objects that must be segmented out of CT or MRI scans semi-automatically. The platform presents a large number of other features and tools for the specific needs of different surgical procedures.

In a research context such as ours, new tools have to be added to the platform on a regular basis. This must be easy to achieve both in terms of programming the new tools and implementing them in the platform.

The Viewing Wand rendering code uses a library called *IAP (Imaging Application Platform)* that was developed by ISG Technologies. Previous work done at our lab using this library has shown that it requires an extensive period of time to learn how to use it, and that it does not allow access to certain steps of the rendering pipeline. Both of these limitations are of importance to our group. Since some student projects must be completed in a year or so, spending a few

months to familiarize with a library that is used in that context only is not efficient. Moreover, the fact that part of the rendering pipeline is a “blackbox” increases the coding time, and sometimes completely prevent certain tasks to be performed.

For these reasons, the IGNS lab at the MNI has created its own visualization platform for clinical applications, *VIPER (Visual Integration Platform for Enhanced Reality)*. It is coded in C++ using two well-known and widely used libraries, namely Motif [8] for the graphical user interface (GUI) and OpenGL [25] for rendering-related tasks. Once the basic structure of VIPER was constructed, implementation of new tools was found to be easier than it was in the Viewing Wand environment. The next section presents the basic structure of our platform.

4.1 Neurological Navigation in VIPER

The first basic imaging feature that was developed on our neuro-navigation platform was the ability to display 2D slices of volumetric MRI scans. Since the data sets that we are interested in usually result from the acquisition of images of a complete 3D volume (volumetric scans), it is possible to view slices from different orientations. In VIPER, the scans can be visualized from the three standard orientations: sagittal, coronal and axial (or transverse) views. We refer to these orientations as *standard tri-planar views*. The user can define which slices of the volume he wants to visualize and can browse continuously through the volume. The way this is done will be explained in detail later. The GUI will be described first.

VIPER appears as a full-screen window divided into five distinctive regions (Fig. 4.1). Along the top, a set of menu buttons allow the user to perform different operations and to select settings, similar to other software applications. A region dedicated to displaying volumetric information (such as intensity values and cursor position in different coordinate systems) is placed just below the menu buttons. In the leftmost region of the main window, a series of icons allows for standard routines to be called quickly (loading up volumes, dropping markers...). Below that, a control panel contains tools for modifying the appearance of displayed images. Finally, the largest part of the window is reserved for the display panel where the images can be visualized. This panel

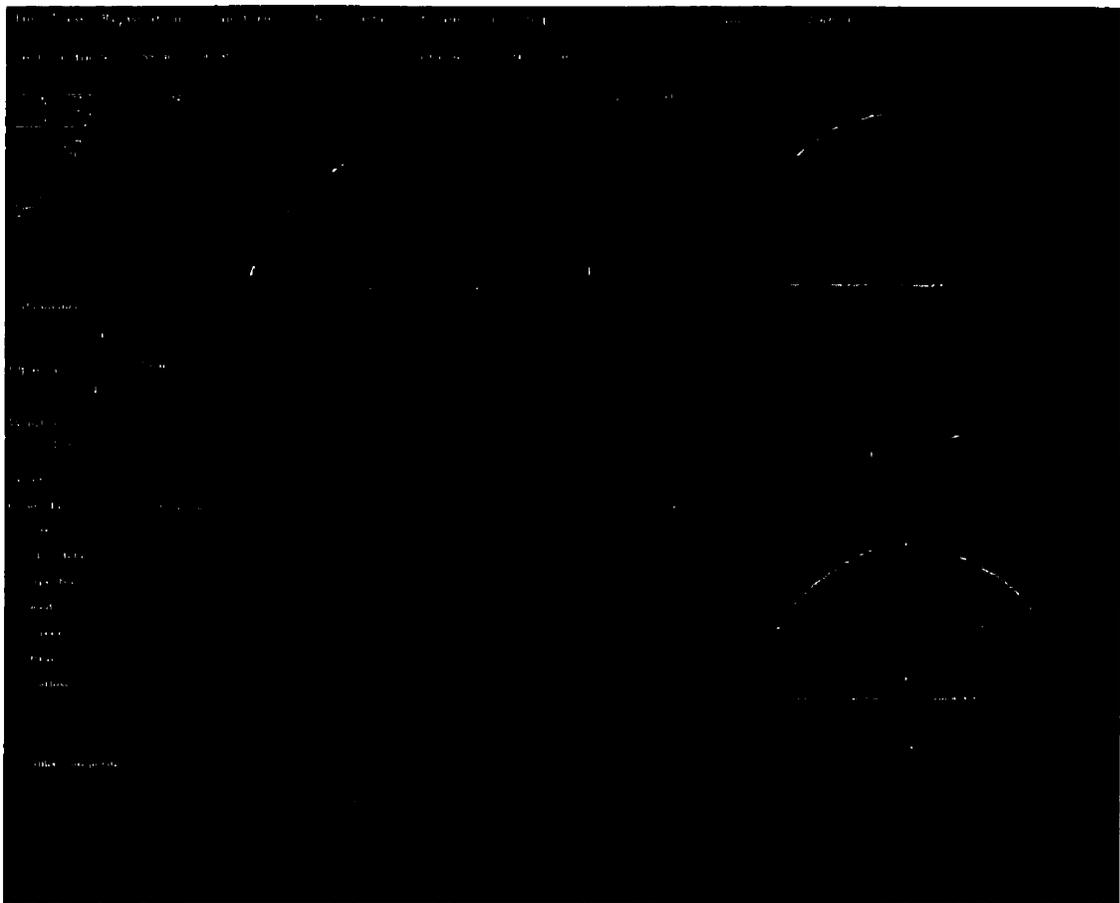


Figure 4.1: The VIPER platform.

is initially divided into four areas of which three are dedicated by default to standard tri-planar views. The way this large display panel is divided depends on what specific feature of the application is used. In stereoscopic viewing mode for example, the whole window is dedicated to the visualization of the 3D object.

The user can choose which slices of the volume are displayed, by a simple mouse click in any of the three standard visualization panels. For example, by clicking on a point in the axial panel, the sagittal and coronal slices that correspond to the lateral and anterior-posterior position of the point are displayed in the two other panels. In this way, the intersection of the three orthogonal planes defines the point of interest. The voxel coordinates and the world coordinates of the point selected by the user are displayed in the volumetric display panel, as well as the intensity value stored in the voxel. The user can also navigate through the volume by moving from one slice to

the next using arrow keys. It is possible to go directly to a point by entering its 3D coordinate as well, using the “Go to point” option of the View menu.

Other standard features, such as zooming of the image, are also available. The zooming can be performed in three ways: by using the scroll bar on the left side of each panel (to zoom or unzoom progressively on the image), by using the quick buttons to go to twice, three times or half the original size of the image, or by dragging the mouse up and down while the control key and the middle mouse button are held down. When an image is zoomed, it is usually not possible to display it in its entirety. The user can scroll through the image using the vertical and horizontal scrollbars found at the bottom and on the right side of each panel.

Displaying an MRI data set as slices in the standard tri-planar view is fundamental to image-guided neurosurgery, but is usually not sufficient. MRI scans provide anatomical information on the patient’s brain, but only to a certain degree of resolution, usually limited to about one millimeter. Moreover, no functional characteristics of the brain are present on an MRI scan. The use of *Positron Emission Tomography (PET)* scans, *Magnetic Resonance Spectroscopic Imaging (MRSI)* scans, *functional MRI (fMRI)*, and other functional or chemical imaging modalities has already been proven useful in the neurosurgical context [33]. In our work, we were mostly concerned with prior anatomical classification (use of atlases) of the brain (see Chapter 3). The next section presents the way these modalities are integrated in the VIPER environment.

Functional imaging modalities usually present very poor spatial resolution, i.e. on the order of 3-6mm, and slice-by-slice visualization presents no anatomical information at all. They appear as diffuse clouds of activity, chemical concentration, etc. They need anatomical context in order to be understandable. This context is provided by an MRI scan or a CT scan (usually approximately 1mm resolution) to which they are properly registered in space [30, 32], so that they can be displayed on top of a structural scan (MRI, CT) at the proper position and orientation.

On the other hand, even MRI and CT scans do not always provide enough anatomical insight into a region of the brain, either because of low contrast or limited resolution. This is the case in thalamotomies where the specific nodes (Vim, Vci...) which must be identified demonstrate no differences in contrast on MRI scans or on CT. Superimposition of precise atlases on top of

these scans is necessary in such cases.

The superimposition (or merging) of many scans in VIPER is achieved according to the philosophy of adding complementary information (functional, chemical, taxonomic...) to a primary structural scan (MRI or CT). VIPER classifies the volumes to be loaded in two categories: one *primary volume* over which up to seven *secondary* volumes can be attached. The primary volume is loaded first, followed by the secondary ones. The secondary volumes are limited in spatial extent to the size of the primary volume. Since the primary volume (usually an MRI scan) provides anatomical context to the other scans, there is no need to display data that falls out of the spatial extent of the primary volume. The control panel on the left part of the main window controls the relative appearance of each individual scan.

Any volumetric scan consists of a list of voxels that are each assigned a given position in space. An intensity value is stored in each of these voxels. This information can be displayed on the screen in numerous ways. The most standard display is a greyscale mapping from intensity values (stored in every individual voxel) to luminance intensity on the computer screen, such that the voxels with the lowest value appear in black while the ones with the highest appear in white. Intermediate voxel values show up on the screen as different shades of grey. The lowest and highest value to be displayed can be chosen using the *Window* slider and *Level* slider in the control panel. The Level slider controls the middle value l of the greyscale (the voxel value that will be displayed in mid-grey) while the Window slider defines the range w of values to be displayed. Voxels with value $l - w/2$ (or less) appear as black while those with values of $l + w/2$ (or more) appear as white.

This method is adequate when displaying MRI scans or CT scans on a computer screen, but does not allow for unambiguous superimposition of other modalities. The use of colours is indicated in this case. Six additional colormaps^a have been added to the standard greyscale, so that a number of different volumetric scans could be superimposed and visualized simultaneously without ambiguity. An opacity slider controls the relative blending of the different images that appear at the same time in a display panel. The scale value of the slider runs from zero to one.

^aSpectral, Hot Metal, Red, Green, Blue and Yellow colormaps.

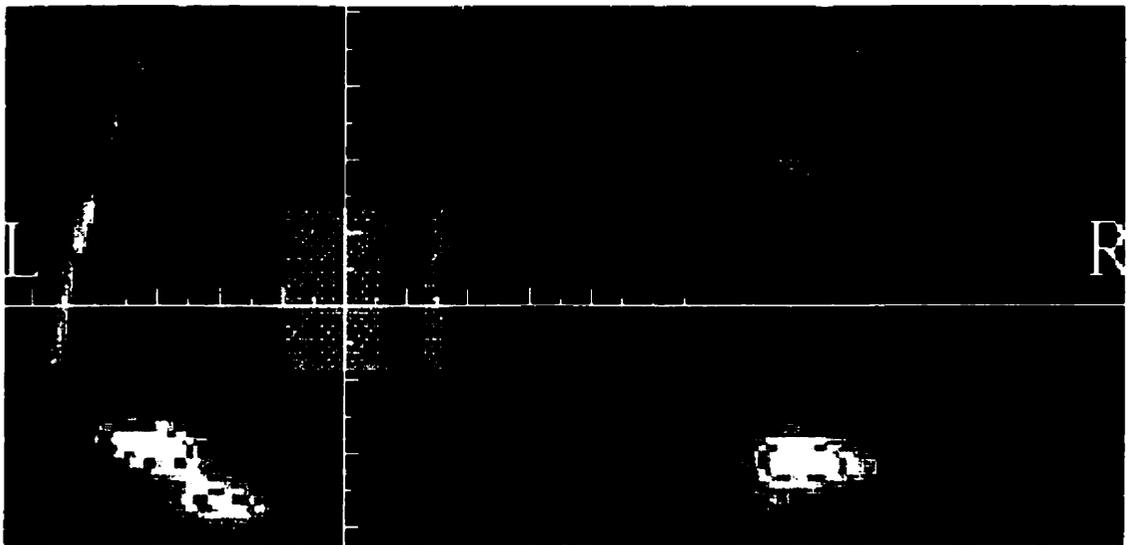


Figure 4.2: Coronal view of a patient's MRI with PET and MRSI scans superimposed

The two oval regions in the bottom of the image represent PET activation from a repetition task. The large rectangular region represents Choline concentration in the tumor.

The primary volume (MRI or CT) is drawn on the screen first, followed by the secondary ones which are "blended" on top of this image one by one. For instance, if both a PET scan and an MRSI scan must be displayed on top of an MRI data set, the PET scan is merged with the MRI according to the PET opacity value, then the MRSI image is merged with the already merged image (of PET and MRI) according to the MRSI opacity value (Fig. 4.2). For this reason, the opacity value of the MRI is always set to one (there is no gain in merging it with the previous image, which is a black screen), except when the user wants to mask it completely by setting the opacity value to zero.

The opacity of the slices as well as the window, level and colormap for every individual scan can be adjusted independently. When modifying the display attributes of one particular scan, its name must first be chosen from a list that appears when the list cascade button is clicked.

4.1.1 Trajectory Views

Tri-planar views often provide enough information on how the anatomical structures are related to one another to appropriately guide the neurosurgeon. However it is possible in certain situations that a cut taken under a non-standard orientation through the volume is necessary. For example, whenever a needle is inserted inside the patient's brain at some given angle away from the standard orientations, the neurosurgeon might want to see where it is passing through the brain. Displaying the needle in standard tri-planar views will make it appear as a small dot on these images, since it would only cross through the slice at one point. However, by cutting the volume in one of the planes where the needle is lying, it is possible to see its complete path in one image, and thus to ensure that no critical structure will be harmed.

The orientation of the cutting plane is determined by calling the "Set angles" function in the View menu. The three angles are defined as *declination*, *azimuth* and *twist*, and they correspond to the well-known Euler angles (Fig. 1.2). The declination gives the angle of rotation around the X -axis (running from left to right). The azimuth defines the rotation around the "new" Y' -axis (Y'), which is simply the Y -axis after it has been rotated around the X -axis by the declination angle. Similarly, the twist angle corresponds to the rotation around the updated Z -axis (Z''); this axis has been rotated first around the X -axis (Z'), and then around the Y' -axis (Z''). It can easily be shown that the rotation matrix corresponding to this transformation is the same as the one that would be obtained by first making a rotation by the twist angle around the Z -axis, then by the azimuth angle around the original Y -axis and finally by the declination angle around the original X -axis.

Two different orientations can be chosen in order to look at slices taken at angles other than the standard tri-planar views. The "gunview" orientation shows the plane perpendicular to a virtual needle that would be inserted in the patient's head with some given declination and azimuthal angles. Mathematically, a plane passing through a point is completely characterized by two angles, namely the declination and azimuthal angles in our case. However, when the slice corresponding to that plane is to be displayed on the screen, an extra degree of freedom appears as the orientation of the 2D picture with respect to the computer screen. The twist angle then

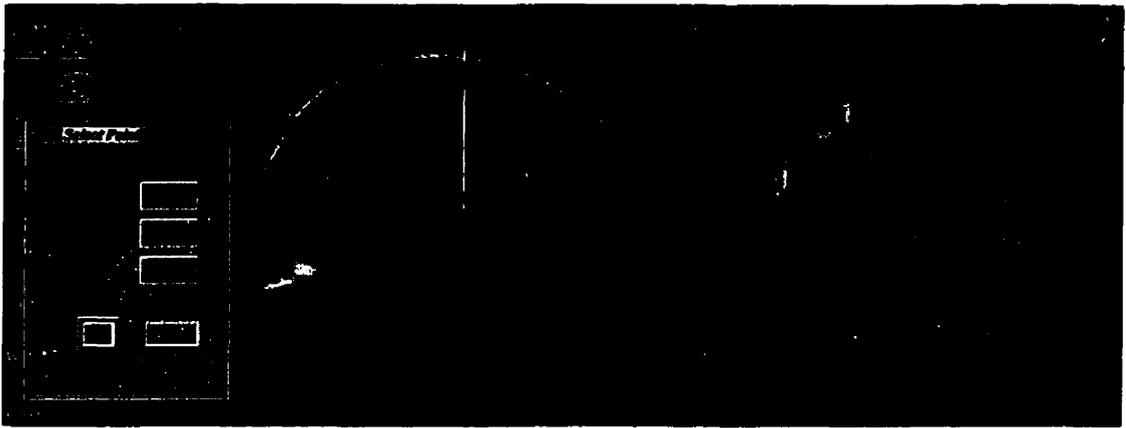


Figure 4.3: The two trajectory view modes in VIPER.

defines the up vector for that plane, i.e. the orientation of the plane itself in the display panel. One can think of the needle being used in the same way as aiming a gun at a target. What the neurosurgeon would see in doing so is the picture that is drawn on the screen, hence the term “gunview” for this view mode (Fig. 4.3).

It is also necessary to be able to look at the volume as if it were cut in the plane where the needle lies. The second oblique view available is called “ultrasound view”, since it provides an image similar to the one obtained by acquiring an ultrasonic image. This time, the declination and azimuth angles define a line that lies in the cutting plane. The plane is not characterized until the twist angle is entered, since the same line lies in different planes, generated by rotating one plane around the line axis. The twist angle defines which of these planes should be chosen as the cutting plane. The orientation of the line (the needle) in that plane defines the up vector, so that the needle always appears to come from the top of the 2D image drawn on the screen (Fig. 4.3).

4.1.2 Markers

Most of the other standard features often found in other volumetric navigation platforms are also available in VIPER. The user can drop markers, or annotations, at any point in the volume. The markers appear as circular spots drawn on top of the 2D images of the slices, with numbers, letters or text next to them. The markers do not appear only on the slice where they were initially

defined, but also in the next and second-to-next slices of their original location, where they appear as smaller circular spots. By navigating continuously through the volume, they look like small spheres being cut off by the slice plane. The markers can be turned on or off, their annotations edited at any time, and distances between two markers can be calculated as well. This last feature is very helpful in obtaining the AC-PC distance, and in giving the distance to be crossed going from the skull to a given point in the brain. The marker list is automatically saved when the user modifies it. Whenever the volume is loaded up again in VIPER, the marker list is reloaded at the same time.

4.2 3D Tools in VIPER

The type of visual information that is available in the 2D display panels is often a limiting factor. The fact that only slices through the volume can be visualized rather than complete 3D structures forces the neurosurgeon to mentally extrapolate the extent of these structures outside of the slice, which is often a difficult task. It is particularly difficult to reconstruct the overall shape, size and orientation of a region of the anatomy or of a lesion site in the brain. Neither is it trivial to assemble different neighboring parts together in space and to imagine how they are related to one another. A neurosurgical tool such as the leukotome is another example of a possible advantage of a 3D-viewing mode instead of 2D. It has to be rotated around an axis which itself was rotated with respect to the standard axes: it is a complicated mental task to follow its position with respect to the critical structures that must be either avoided or excised. In order to compensate for these limitations found in 2D visualization, we have developed 3D imaging tools in VIPER that allow for more insight into the global geometry of objects presenting some neurosurgical interest.

4.2.1 Surface Rendering

There are many ways of representing 3D data. We are interested in the following: sets of points, meshes, surface rendering and volume rendering. We show later that points and meshes are only

intermediate steps of the surface rendering pipeline, so that we consider only two general processes in performing 3D rendering, namely surface rendering and volume rendering. Surface rendered images come from a manual or semi-automatic classification (or extraction) of surfaces from a volumetric data set. Series of points that define the boundaries of the surface object to be rendered must first be extracted. The computer renders the objects by calculating the view that would be seen if the surface was illuminated by spotlights reflecting on the surface of the model object.

In order to be able to perform the model calculation for light reflection, the computer needs a list of individual polygons with a complete description of their position, size and orientation in space. By rendering polygons differently depending on their position and orientation, the rendering algorithm projects to the screen an image that corresponds reasonably to what an observer would see if a real 3D surface were seen on a photograph. Hence a series of points found on the surface to be modeled must be joined together to create a list of polygons. To display 3D information coming from an MRI scan using surface rendering, segmentation of the different structures can be performed semi-automatically using a “filling” (or seed-growing) algorithm that searches for intensity values within a certain range starting from a given point in 3D. Another approach involves manually drawing a “labelling” on top of the region to be surface-rendered. At our lab, both of these operations can be performed using the “Display” software developed by D. MacDonald [28].

This method works well when one wants to create a surface from the boundaries of a tumour or to outline the ventricles in the brain, since these regions show significant contrast with surrounding tissue. However in our case most of the thalamic nodes show no difference at all in MRI intensity value with respect to one another, making any kind of segmentation based on the MRI data set impossible. We faced the same problem earlier in the 2D case since the regions of interest (the thalamic nodes) could not be distinguished on the 2D slices of the MRI data set. Our 3D surfaces describing the anatomy of the thalamus can still be created using the segmentation algorithm on the atlas itself instead of on the MRI data set. Since the volumetric atlas is composed of 3D regions having been assigned discrete intensity values, the boundaries between

regions are sharply defined. Thus the boundary-finding algorithm works very well in this case.

Since the information available on the atlas is more boundary-like than that on the MRI data set (where the intensity value of each voxel contains theoretically valuable data), the surface segmentation and surface rendering method is well adapted in the task of displaying that information in 3D. The difference between the atlas and the MRI data set comes from the fact that the uniformity of intensity values in one given region of the atlas adds no information once the geometric boundary of that region is found, while a volumetric region on a MRI data set contains some variation in intensity within its boundary. Another method of 3D rendering that is better adapted to this kind of data set will be presented in section 4.2.2.

In VIPER, up to ten different surface objects can be loaded and visualized simultaneously. We have limited this number to ten for speed considerations, since the 3D viewing panel loses interactivity when the number of object is too large. The CPU time-cost of rendering surfaces is significant due to the large number of polygon lighting calculations. Surfaces are loaded by pressing the “3D-loading” button on the quick button toolbar. Once at least one object is loaded, a “3D-attribute” pop-up window appears on the screen (Fig. 4.4). This window contains the control panel for the rendering of 3D objects, and is similar in many ways to the static control panel for the 2D viewing panels. The object whose display attributes are to be modified must first be chosen from the list appearing when the list cascade button is clicked. A scale scrollbar controls the opacity of each individual object in the same way as the 2D slices were superimposed. An opacity value of one means that the object is completely opaque, while a value of zero makes the object invisible. Instead of having to choose between different colormaps, the user indicates what color he wants the surface to appear from a choice of seven colors.

The 3D window always displays the three standard axes in red, yellow and blue for the X , Y and Z -axes respectively. The position of these three lines in 3D is given by the 3D-coordinate of the current position of the crosshairs in the standard tri-planar views. The surfaces originally appear with an opacity value of one, i.e. completely opaque. When all the objects are opaque, they are drawn one after the other with their depth values being stored in the depth buffer.

The depth value, also called z -value, is defined for every single polygon and corresponds to

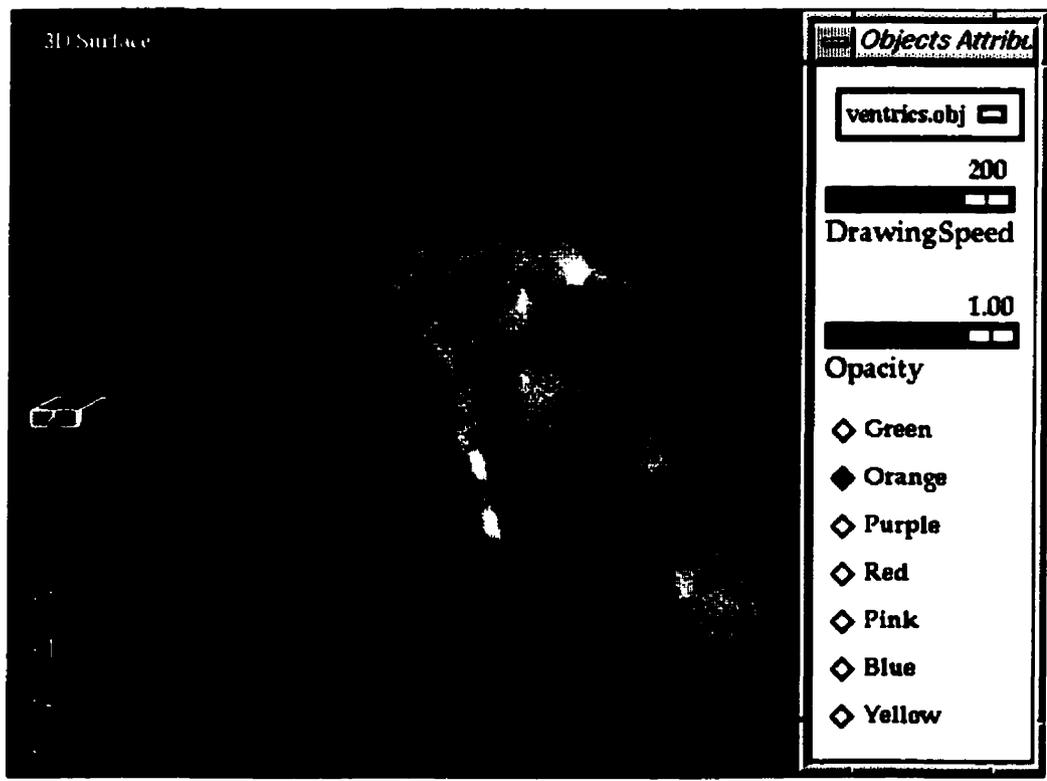


Figure 4.4: Surface rendering of the Superbrain's ventricles in VIPER.

the modeled distance from one polygon to the position of the observer in 3D space. This information is necessary to ensure that the algorithm does not draw an object that should be hidden by the other just because it is drawn last. The problem is more complicated when translucent and opaque surfaces must be rendered at the same time. For instance if a translucent surface is drawn first and its depth values are stored in the depth buffer, another surface getting drawn afterward would not appear on the screen if it were placed at a greater distance from the observer than the previous translucent surface. This is obviously not the behavior that is expected from a translucent surface. The idea is to first draw all the opaque objects, and then the translucent ones while making the depth buffer “read-only”. The depth test still verifies that the translucent surface should not be drawn if it is hiding behind an opaque object, but does not modify the depth buffer even if it lies in front of other objects. In that way, another translucent surface that should be drawn between the closest translucent surface and any other objects behind it will still be drawn since the depth buffer does not “know” that the first translucent surface is indeed closer.

Perhaps the greatest advantage of a 3D viewing panel is that the volume can be inspected from different points of view. The user should then be able to rotate the volume (or to rotate his point of view, which is exactly the same thing in our case), as well as to translate and zoom it. The translation is done in the 2D panels, since the 3D surfaces are centered accordingly to the 3D coordinate of the crosshairs in the 2D views. The rotation and the zooming are performed via the mouse, in a simple and intuitive way. By moving the mouse up and down while holding down the left mouse button and the control key, the observer’s point of view becomes closer to or farther from the scene.

Rotations are more complicated to model. In principle, three parameters need to be specified to completely describe a rotation. However, if at any time the rotation increments are defined around the three standard axes with respect to the observer, i.e. with the Z -axis facing toward the user (coming out of the screen) and the X and Y axis in the plane of the screen, the rotation about this Z -axis usually does not need to be controllable by the user. Keeping in mind that rotations in this context are useful to show the object from different points of view (revealing regions that

might be hidden), performing a rotation around the Z -axis (which can be seen as a rotation of the plane of the screen) does not help in giving better insight on the geometry of the objects since the same 2D-image would simply be rotated around its centre. It seems straightforward to use the movement of the mouse cursor in the screen to increase or decrease the value of the rotation angle around both the X -axis and the Y -axis. Left multiplication of the “increase rotation” matrix by the previous rotation matrix (which in turn resulted from the accumulation of prior rotations) should result in a new rotation matrix that will show the object from a new point of view. The problem with that method is that it accumulates round-off errors at every small increment of rotations. These errors modify the rotation matrix in such a way that after a while, the matrix cannot be factorized into rotations only. When these errors accumulate significantly, the matrix might demonstrate a “projection” behavior; i.e. one axis can be completely folded onto another, making rotation around this axis impossible. This behavior is known as a *gimbal lock* [25].

In order to avoid that problem, another parametrization of the rotation is used. This method is called the *quaternion* representation. Quaternions are 4D normalized vector objects representing a rotation. The normalization of the quaternions ensures that round-off errors cannot accumulate. The movement of the mouse in the 3D window is translated in a small rotation represented as a quaternion, and then this quaternion is added to the previous quaternion to form the new representation of the rotation. This last quaternion is finally normalized and translated back into a rotation matrix that is applied to the objects prior to the rendering steps. The user controls the rotation in the 3D window by moving the mouse in any direction while holding down the middle mouse button and the control key.

When the total number of polygons to be drawn and rendered is greater than ten thousand or so, the updating of the surfaces starts to be somewhat sluggish, i.e., the rotation and zooming appears more as frame-by-frame images than as a continuous movement of the objects. This is true for the computer on which the platform was developed, an O2 SGI unit. In order to avoid such behavior, the rendering algorithm automatically changes to a quick drawing mode as soon as the user tries to zoom or rotate the object. By default, this quick mode consists of rendering only one polygon out of twelve. This parameter is controllable by the “Drawing Speed” scrollbar

in the 3D-attributes pop-up window. The value of this parameter can vary from 1 (every polygon is then drawn) to 256 (one polygon out of 256 is drawn).

4.2.2 Volume Rendering

Surface rendering is well indicated for 3D representation of pre-defined boundaries between regions, such as atlas regions describing thalamic nodes. It should be stated that the restrictive representation of the thalamus into distinct regions does not reflect exactly the functionality and anatomy of the thalamus, but these are the limitations of an atlas-oriented model. Once this model is chosen, surface-rendered objects constitute the best way to represent its content. However, even though the MRI volumetric data set does not present enough information relating to the thalamic nodes to adequately guide the neurosurgeon, it nevertheless complements the atlas information by offering spatial and general anatomical context.

We have already discussed the need to have underlying MRI slices to provide spatial context to the atlas in a 2D representation. The same problem arises in 3D. Some anatomical references that are familiar to the neurosurgeon should be provided to enhance the surface representation of the thalamic nodes. Since the MRI data set is volumetric, in the sense that voxels contain (in principle) independent information from one another, no definite boundaries are present *a priori*. Of course, some structures such as the skull or the ventricles are clearly outlined, but a surface segmentation of these structures is not necessarily the representation that will look the most familiar to the neurosurgeon. One of the most important characteristics of the brain is its white-grey matter differentiation, and the complex connectivity of regions of white or grey matter (and the fact that there are many regions of mixed white-grey matter) does not allow for a simple segmentation.

The volume rendering technique is a representation that keeps all volumetric information available in the data set. The construction of a volume-rendered 3D image consists of a set of virtual rays cast through the volume onto the observer's position. It is similar to a complete set of X-rays taken from all possible angles around the volume to be investigated. The algorithm (in its simplest form) operates in the following manner: the ray starts from the opposite site of

the data set (with respect to the observer's point of view), then travels through the volume. As it strikes a voxel with some non-zero intensity value, this value is added up to the previous sum of other voxels seen by the ray, and so on until the ray reaches the observer's eye. When the user rotates the object, a new set of rays are cast through the volume under a new orientation. This operation is intensive in terms of computer time, and is the principal drawback of volume rendering.

On the other hand, volume rendering is a very versatile 3D-visualization technique. The way the virtual rays can accumulate volumetric information when passing through the volume is not limited to simply adding up voxel values. A well-known example is the *Maximum Intensity Projection (MIP)* technique essentially used for the 3D rendering of angiograms. Instead of adding up the voxel values, the algorithm stores the maximum value it encountered on its path. On an angiogram, this results in showing only the largest vein or artery found along a particular line of sight.

The image that is projected to the screen when the voxel values are simply added along rays has the appearance of a translucent volumetric data set. The ray can get through, but the final value that goes to the screen depends linearly on the voxel intensities traversed by the rays. It is possible to weight the voxels with respect to either their positions or their intensities. For example, by assigning more weight to the voxels that are found closer to the observer, a more realistic scene is rendered. The deeper a structure is located in the patient's head, the less effect it will have on the image, so that the picture mimics a partially translucent object. Some specific range of intensity values can also be more heavily weighted, in order to enhance a characteristic of the tissue in place. For instance, by knowing in what intensity range the white matter lies, it is possible to emphasize these voxels so that the white matter will appear brighter than the rest (Fig. 4.5 and Fig. 4.6). This is quite similar to the semi-automatic segmentation method, except that volume rendering allows for "fuzzy" regions to be displayed as they are, while the surface rendering method defines in a binary manner whether a point is on one side of a boundary or the other.

Displaying a thalamic atlas in 3D using surface rendering shows precise definition of the



Figure 4.5: White-matter weighted volume rendering of the Superbrain.



Figure 4.6: Grey-matter weighted volume rendering of the Superbrain.

nodes, but gives little context as to where these structures are localized in the brain. When the user starts rotating the surfaces in the 3D window, the orientation of the nodes with respect to the surgical space can be lost, so that the neurosurgeon loses track of the direction from which he is looking at the atlas. By adding volume rendering of the MRI data set (or a part of it) at the same time, the surgeon obtains the global context he needs to understand the information available from the atlas.

4.3 Neurosurgical Tools

During a thalamotomy, two neurosurgical tools (the electrical stimulator and the leukotome) are inserted deep within the patient's brain. Since no craniotomy is performed during this operation, the neurosurgeon never sees directly what is happening with the tools inside the brain. The only proof that the tools are correctly in place is found on the many x-rays that are taken during the operation (see Chapter 2). The x-rays only show a transmission image through the brain, where only the ventricles appear clearly. The neurosurgeon must mentally reconstruct both the geometry of the anatomy *and* the position of the neurosurgical tools with respect to that anatomy.

In VIPER, the geometry of the anatomy is already displayed either on 2D slices or in the 3D window as surface-rendered objects. It is also possible to display the position and geometry of the neurosurgical tools with respect to the thalamic nodes, on the 2D slices as well as in the 3D window.

The two tools were first optically scanned on a *HP ScanJet IIcx*. A translucent acetate sheet with millimetric grid drawn in black was first put on the scanner window before the tools were aligned with the millimetric scale and scanned (Fig. 2.2). These pictures were then zoomed so that the geometric extent of the tools could be measured. For the stimulator, only the position of the electrode tip with respect to the hole in the shaft was measured, since the specific geometry of the wire itself is of no surgical interest. The position of the tip was measured for all fourteen possible settings (or extents) of the electrode tip.

For the leukotome, however, a complete description of the geometry of the metallic wire was

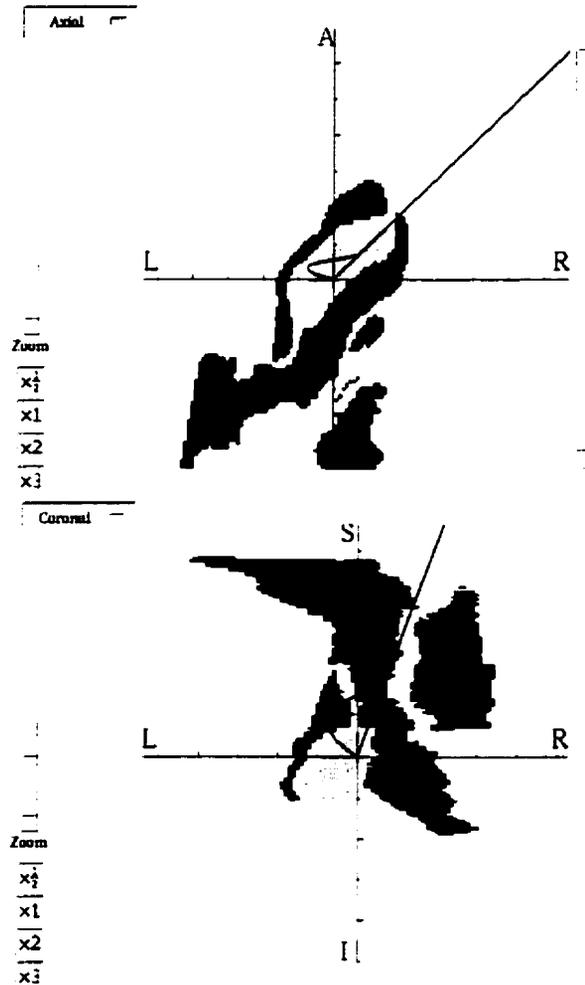


Figure 4.7: Example of the orthogonal projection ambiguity.

The leukotome loop position on the axial view (top) can mislead the surgeon in thinking that its extent is limited to the Gpm (light grey), while the loop is also placed in the internal capsule and external pallidus (dark grey), as seen on the coronal view (bottom).

needed in order to model it appropriately. Samples of the distance between the metallic wire and the shaft of the leukotome were taken every millimeter on all ten millimeters over which the wire can protrude from the shaft. This was done for all seven different settings (protusions) of the leukotome loop. Those values were entered into the software, so that the neurosurgical tools could be drawn in the 2D and 3D windows.

The position of the leukotome with respect to the co-registered atlas and MRI data set is fully controllable by the user. The position of the target point (the point where the tip of the leukotome should be) and the two angles of entry (declination and azimuth) must be specified first. This position can be given either in the native MRI coordinate system, or more commonly in the stereotactic frame coordinate system. The three settings of the tool must be typed in, namely the twist angle (rotation of the tool around its own axis), the offset and the microscrew setting. The offset corresponds to a translation of the tool along its own axis. The microscrew setting defines the extension of the tool from the shaft. All of these parameters are entered in the pop-up window that appears when the Thalamotomy button is selected on the VIPER menu bar.

The neurosurgical tools are usually inserted into the patient's head at an arbitrary orientation with respect to the three standard axes; i.e. the shaft of the tool does not lie on one of the three standard planes (coronal, axial, or sagittal). This means that when one of the tools is within the patient's brain, a slice taken in one of the three standard planes would only show one small dot where the shaft crosses the plane, and possibly a second one (even smaller) if the plane runs across the wire. This would not give great insight on the position of the tools with respect to the anatomy. If an orthogonal projection of the complete tool is drawn instead of showing only the dots where the tools cross a given plane, the position of the tools can be better understood. The main drawback of this method is that it could mislead the neurosurgeon into believing that the tool really lies in one of the three standard planes when it does not. This would cause a problem when using the leukotome since the projection of one part of the thin metallic wire can be projected onto an anatomical structure (in the 2D view) while it is not really touching it (Fig. 4.7).

Once all the parameters are set appropriately, which slices in each of the three standard planes should be displayed? We know for sure that only one or maybe two points belonging to the

tool will be found on the chosen plane. The interface works somewhat differently depending on whether the leukotome or the stimulator needs to be displayed. For the stimulator, the most important point is not on the shaft but rather at the tip of the electrode. Hence the slices to be displayed are chosen so that all three contain the tip of the stimulator electrode. The interface updates the changes to the views in exactly the same way as if the user was calling the “Go to point” function with the position of the tip of the electrode as the 3D coordinate. As a matter of fact, the shaft and wire of the stimulator are drawn to the screen to give an idea of how the stimulator is oriented in the patient’s brain, but this information is not essential to the procedure. It would be sufficient to display only the position of the tip of the stimulator electrode with respect to the anatomical structures modeled by the atlas.

For the leukotome, the situation is more complicated since no single point on the metallic wire is more important than the others. We decided somewhat arbitrarily that this point should be the closest to the tip of the shaft where the metallic wire protrudes. This choice leads us to the projection ambiguity problem discussed previously (Fig. 4.7). This problem is avoided by investigating the position of the leukotome with respect to the anatomy on a trajectory slice in which the leukotome lies. By doing so, the correspondence between the metallic wire and the atlas is not subject to the projection inconsistency. However, the standard tri-planar views can still be useful in providing general anatomical context (Fig. 4.8), even though the neurosurgeon must be aware that the leukotome position is only an orthographic projection onto the displayed slice.

During the surgical procedure, the excision is performed by rotating the leukotome around its shaft axis. The neurosurgeon must be able to determine what the leukotome path through the anatomy will be when performing the excision. This is difficult to show in the 2D panels, since rotational rather than translational movement of the leukotome implies that it must be going out of the plane in which it was originally lying. This is one of the main reasons for the implementation of the neurosurgical tool models in the 3D window. The tools are *surface-rendered* using long cylinders as the surface model for the shafts, and series of small contiguous protruding cylinders representing the wires. The positioning of the neurosurgical tools is performed using

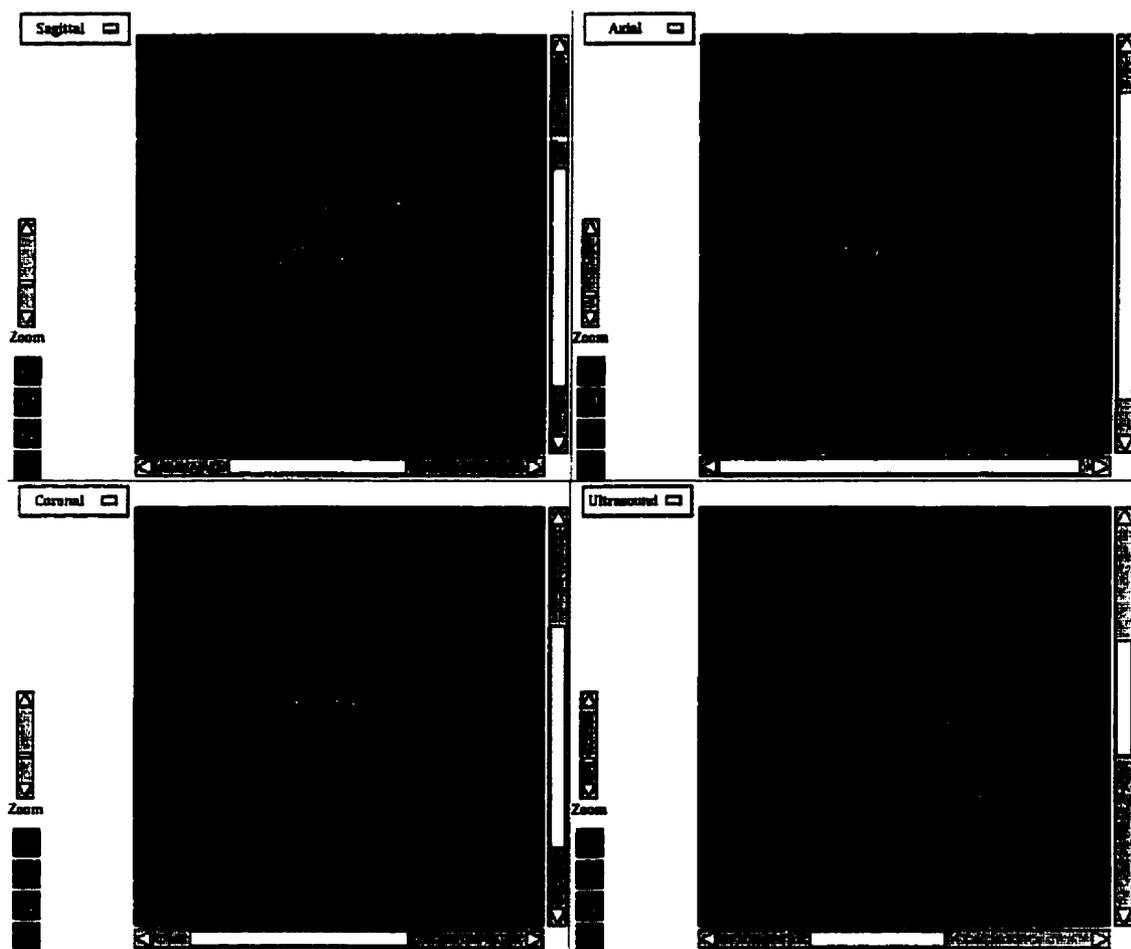


Figure 4.8: Standard tri-planar views and trajectory view of the leukotome model, the atlas and an MRI data set.



Figure 4.9: The 3D leukotome model inside the Vim node.

the same pop-up window as for the 2D panels, ensuring that the 3D view is consistent with the 2D views. The 3D window also offers a more interactive way of modifying the settings of the tools. By pressing the 'k' and 'l' keys, the user can continuously rotate the leukotome model clockwise or counterclockwise respectively. By pressing the 'i' or 'o' keys, the modeled wire can be taken in and out of the shaft. The neurosurgeon can then see exactly how the leukotome penetrates the region of excision as it is rotated around its shaft (Fig. 4.9).

4.4 Stereoscopic Viewing

In the 3D window we try to present structural features such as depth, geometry and the relative position of objects. These features cannot be seen on two-dimensional slices. However, the computer screen is still a two-dimensional plane, and we are limited to showing projections of 3D objects onto that plane. The user gets more insight into the 3D geometry of the objects by interactively moving the object in space, but is still limited to viewing a series of 2D pictures.

The strongest depth-perception cue is called occlusion [43], and consists in the masking of

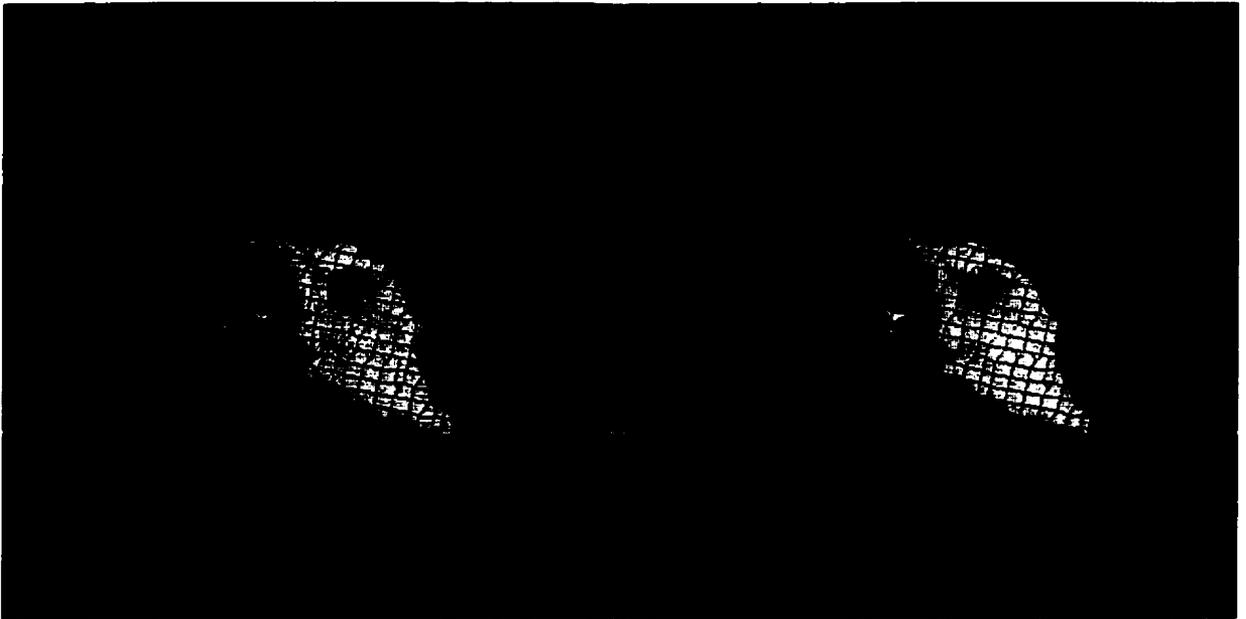


Figure 4.10: Stereoscopic image pair showing a mesh model of the cortex arranged for cross-eyed viewing.

an object by another. This cue is already present in the VIPER 3D-display environment (see earlier discussion on the depth buffer, section 4.2.1). The next strongest cue, stereopsis, relies specifically on human binocular vision. When an observer looks at an object with two eyes, the images projected on the retina of the left and right eyes are slightly different. The brain is able to use this slight disparity to retrieve information about the position of the object with respect to the observer. By presenting a pair of 2D pictures corresponding to the two slightly different views that would be obtained for each eye when looking at a 3D scene, an observer has the impression that he is actually looking at the real 3D scene rather than to 2D projections of it (Fig. 4.10).

There is a way to present the two disparate views of a 3D object to the user on a computer screen. By having the user wear a pair of glasses that allows its left and right lenses to be shuttered alternatively at a known rate, the computer screen can present alternatively the right and left view on the screen at the same rate. Stereographics Corp (San Raphael, California) offers a commercial package (CrystalEyes) for stereoscopic viewing on SGI workstations, consisting of a pair of stereoscopic glasses (the shuttering mechanism relies on liquid crystal technology) with an infrared emitter connected to the computer. The emitter sends a signal to the glasses in

order to synchronize them with the refresh rate of the screen [12, 14, 13, 34].

We have implemented a fully interactive stereoscopic viewing mode in VIPER. An icon button showing a pair of stereoscopic glasses appears on the left side of the drawing area in the 3D panel. By a simple click on that button, the application changes to stereoscopic mode. Of course, the user must wear the Crystal Eyes stereoscopic glasses to see the 3D view. Since the rest of the screen presents only two-dimensional information, it is meaningless to show it in the stereoscopic mode. For this reason, the 3D window is enlarged to fit the complete screen when in stereoscopic mode. Nevertheless the program may be switched back to standard mode at will by relicking on the same button.

In the stereoscopic mode, the user can still interact with the 3D objects, through rotation or zooming. The control panel for the leukotome is also available, so that it may still be rotated, “offset”, or have its wire adjusted.

Chapter 5

Lesion Modeling

This chapter describes some computer tools that were developed to facilitate the surgeon's work in performing the excision step (for thalamotomy and pallidotomy) through modeling of the excised volume.

5.1 Rationale

As stated earlier, the potential advantage of using the leukotome instead of other tools is in its ability to perform a precise tailoring of the region to be excised. Nevertheless, this advantage is gained only if the neurosurgeon is able to use the leukotome with a precision that is at least as accurate as other methods. This means that the two steps involved in the excision, targeting and excising, must be achieved with enough precision so that the tailoring feature of the leukotome can be effectively used.

The targeting step has already been presented in Chapter 2. It is not intimately related to the use of a particular tool, except for the fact that one can choose to get closer to a structure to be avoided when using the leukotome. This is possible since the neurosurgeon can perform asymmetrical lesions with the leukotome, enabling him place it very close to the boundary between two regions and use it to excise only on one side of a boundary. As a rule of thumb, the precision

of the position of the target point should be on the order of half the size of the target volume's smallest dimension. If this is not the case, the tailoring feature of the leukotome is completely outweighed by the imprecision of the targeting. If the neurosurgeon does not know in which part of the target volume he is placing the tip of the leukotome, it is impossible for him to decide which quadrant he should open the leukotome into, and to what extent he should do so.

The target point itself is always set prior to the operation. In principle it is never modified during the procedure, though it could possibly be if the stimulation stage (section 2.2) shows that the pre-operative targeting was inaccurate. Even if the target point is not exactly where it was thought to be, the stimulation stage brings to the neurosurgeon the information he needs to correct the excising protocol. At this point, the main difficulty is to determine how to use the leukotome in order to respect the results obtained from the stimulation.

As stated before, the leukotome is inserted into the head of the patient at some arbitrary angle which can be up to 40 degrees of declination and 15 degrees of azimuth with respect to the frame. Most of the time the frame itself is nearly aligned with the AC-PC plane. Anatomical training of the neurosurgeon is usually stronger in the standard tri-planar views, since atlas books, MRI scans and X-rays are usually presented in these views. This difficulty of "seeing" the leukotome position and orientation in 3D has already been addressed in the previous chapter. Furthermore it remains difficult to decide what is the best excision protocol in some particular cases, even if a good 3D view of the atlas and the leukotome is provided.

We have developed an algorithm that suggests a "good" protocol to use, given a target lesion volume and the position of the leukotome shaft once the target point has been chosen. A "good" protocol is one that minimizes unnecessary excision while maximizing removal of the target tissues. The algorithm does not find the "best" protocol possible, but rather presents a first approximation on what the protocol to be used should be. No perfect solution to this problem exists. In this context, a perfect solution means one where complete removal of the target tissues is achieved while no other tissue is excised. This impossibility rises from the fact that the leukotome wire can only be extended in seven discrete settings, and the number of quadrants is usually limited to sixteen for practical reasons. It is also sometimes topologically impossible

to perfectly fit a given volume by rotating a tool of any shape, e.g., when the volume is donut-shaped. From a geometrical point of view, this means that the leukotome settings do not span the space of possible target volumes.

It might be supposed that the “best” solution would be one which optimizes some function of the two parameters discussed above: by minimizing unnecessary excision while maximizing removal of the target tissues. An additional difficulty arises when taking into account that excisions are absolutely forbidden in some neighboring structures of the target volume (the internal capsule, for example). It then becomes practically impossible to weight the “forbidden-ness” of a region and to account for this in an automatic protocol algorithm. It would be possible to supply some information on the surroundings of the target and to modify the choosing algorithm with respect to that, but the discrimination for some tissues is rather arbitrary and depends strongly upon the kind of response that the neurosurgeon receives from the stimulation stage. Therefore it appears that the algorithm that we developed is more useful if its task is limited to a first approximation of the protocol rather than expanded to take care of a very large list of parameters. The neurosurgeon can not assume that list to be all encompassing.

5.2 Algorithm

The algorithm operates in the following manner. Given a target point, a set of two angles (declination and azimuth of the leukotome) and a target volume, it performs a boundary search and suggests a protocol as a series of leukotome settings for all sectors. The number of sectors can be set by the user, but usually 8 or 16 sectors are used.

The boundary search is performed in ten planes perpendicular to the leukotome shaft, the first being the one crossing at the tip of the shaft. The remaining nine planes are uniformly spaced by one millimeter. When the thin wire of the leukotome is extended by a certain amount, the planes intersect the wire at different distances from the shaft, since the shape of the wire is not rectangular (Fig. 5.1). In every plane perpendicular to the leukotome shaft, the algorithm searches for the average distance from the boundaries of the target to the axis of the leukotome in each sector.

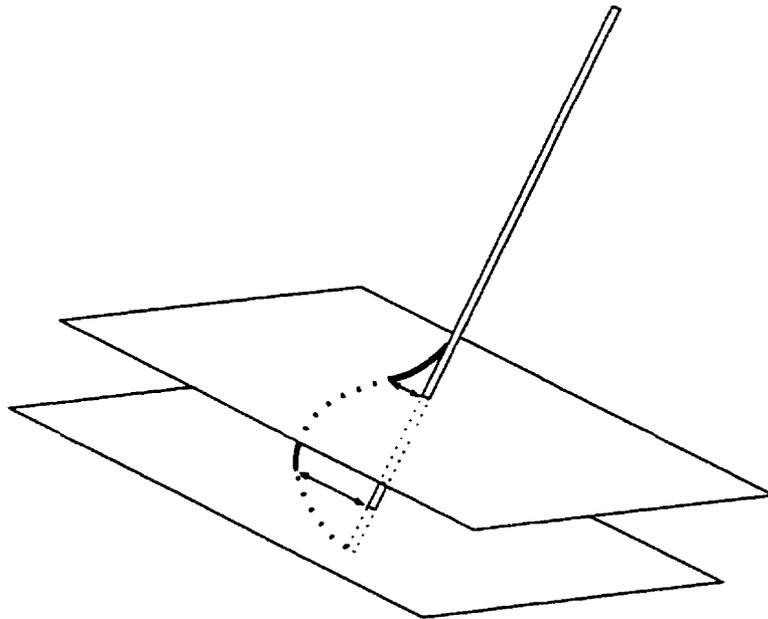


Figure 5.1: The leukotome loop intersecting two different planes at two different distances from the leukotome shaft.

The average distance for one sector is obtained with the following formula:

$$\text{average distance} = \frac{1}{4}(v_1 + 2v_2 + v_3) \quad (5.1)$$

where v_1 and v_3 are the boundaries at the two limits of the sector and v_2 is the boundary in the middle of the sector. It is preferable to do it this way instead of using only the middle value; this prevents making the calculation with only one point that may not be representative of the complete sector (Fig. 5.2). One could say that this method is analogous to a trivial low-pass filter. For each sector, the averaged boundary value is compared with the seven possible settings of the leukotome at the current level. The setting that best fits the average boundary value at that level is stored. This is repeated for all sectors in all different levels. For one given sector, the best settings in all planes (or levels) are averaged. This average setting constitutes the output of the program for that sector.

The program also provides some volumetric information on the validity of the proposed protocol, such as the volumes of the “bad”, “good” and “missed” excision (in mm^3). A “good” excision volume is simply the part of the target volume that will be excised while “bad” refers to

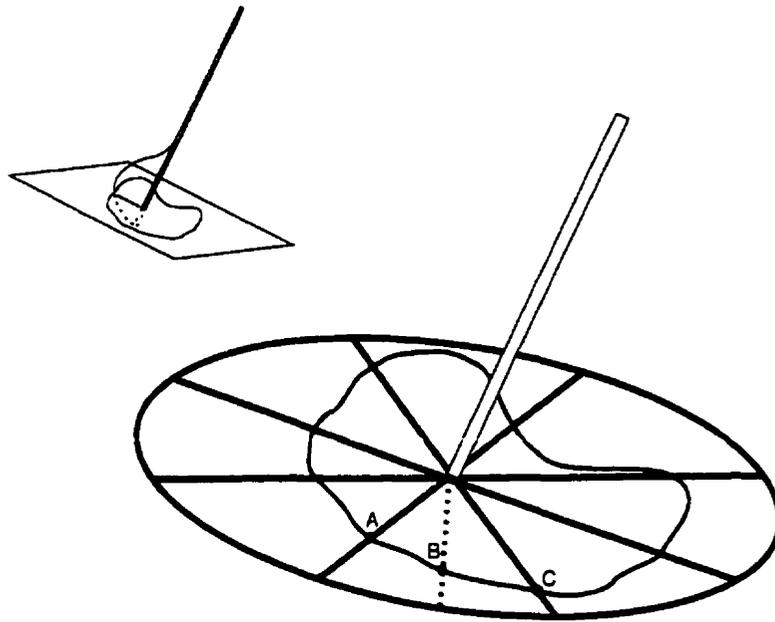


Figure 5.2: The boundary-finding algorithm in one plane. The distances separating point A and C from the shaft represent v_1 and v_3 , respectively. The distance between point B and the shaft represents v_2 .

the non-target volume that is excised. “Missed” volume is the portion of the target volume that would not be excised if the suggested protocol is followed. Once again, this information is of limited reliability since even a very small “bad” excision might have to be avoided depending on the specific region where it lies.

The algorithm also uses a weighting of the planes where the most central levels are given more importance than the others. In these planes the distance of the thin wire with respect to the axis is larger than on the side planes, so that the mean error is reduced.

The algorithm was first developed to be used intra-operatively, in such a way that for each patient it would find a first approximation to the protocol to be used following the stimulation run. Since the target volume is already known before the patient arrives in the operating room and since the target point is defined beforehand as well, it was decided that this tool would be more useful pre-operatively.

From patient to patient, the protocol for the procedure is usually similar. This is due to the

fact that variability of deep brain structures from patient to patient is reasonably limited. The neurosurgeon also tries to enter the brain from similar angles from case to case. The stimulation phase is undertaken to provide the surgeon with additional information on the leukotome position with respect to the anatomy, as well as information on the functionality of the patient's deep brain structures. While this additional information might result in small changes to the standard protocol, the global appearance of the excision map remains the same.

For these reasons, it is likely that the use of our algorithm will be defined as a standard protocol optimizer, even though one can easily run it again in cases where the tool enters the brain at an unusual angle, or when the anatomy is significantly different from average. Rather than a restriction, this is seen as a simplification of the use of the algorithm. It is used once for standardization of the protocol, and used again only when divergence from the general case requires it.

5.3 Virtual Lesion Generator

We have described in the previous chapter how the neurosurgical tools involved in thalamotomy and paladotomy can be visualized in a number of ways (projected on standard tri-planar views, on trajectory view or even in 3D-stereo). Even if this provides a complete knowledge of where the leukotome is positioned with respect to the anatomy at one time, it is still hard for the neurosurgeon to generate mentally the shape and bounds of the total volume that will be excised using one given protocol, and to determine how well it corresponds globally to the target volume to be excised. We have developed a small graphical interface that allows the neurosurgeon to define, in a very user-friendly manner, the protocol he wants to use and to create from it a "virtual" 3D lesion.

This virtual lesion can then be visualized on the VIPER platform either as a tomographic volume (standard tri-planar views and trajectory views) or as a 3D surface-rendered object (Figs. 5.3 and 5.4). It can thus be compared with the atlas to see if any part of it overlaps with structures that should not be excised, such as the internal capsule. The 3D view gives more insight on the



Figure 5.3: A virtual lesion (white) superimposed on the atlas in 2D viewing mode.

global extent of the lesion and on the shape of the lesion with respect to the target. By showing the leukotome in the 3D view at the same time, the user can rotate the tool to a region (sector) where the lesion extent is inappropriate, find a better setting for that sector and then create another virtual lesion. This is significantly faster than a complete leukotome inspection of the excision in every sector. Moreover, the virtual lesion generator appears to be a necessary feature, considering that the neurosurgeon does not rely solely on the 3D view only but also on standard tri-planar views where the inspection must be performed on a slice by slice basis in all three projections.

The graphical interface is presented in the familiar shape of a dartboard divided in sixteen sectors (Fig. 5.5). Each of these sectors is divided into seven concentric parts, corresponding to the seven possible settings available on the leukotome used at the MNI. A total of 112 (7x16) small regions are drawn. Since the protocol usually involves opening of the leukotome at different levels, i.e., with different offsets of the leukotome with respect to the target point, the user can define up to three different levels in which the resection is to be made. By clicking in a region, the user highlights the corresponding sectors from the middle of the dartboard up to the

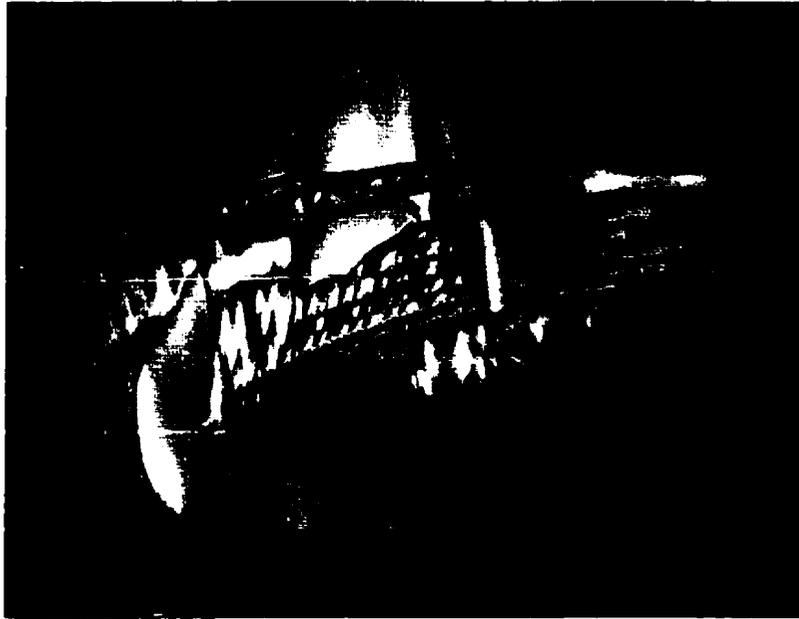


Figure 5.4: A 3D view of a virtual lesion (flat grey) of the Gpm (textured grey).

selected region. By clicking again in the same region, the user toggles only that small region. By selecting the same sector but in a more central region, only the regions inside the first selected are highlighted, while the others are set to black. To clear a sector, the user merely clicks the mouse in that sector, but outside the largest ring. Menu options allow the operator to clear a complete level or to clear all three levels at once. In the bottom part of the window, the user sets target coordinates in frame space as well as the declination and azimuth angles of entry. The user goes from one level to the other by selecting any of the three toggle buttons in the bottom part of the window. The highlighting appears in different colors from one level to the other to clearly differentiate between them.

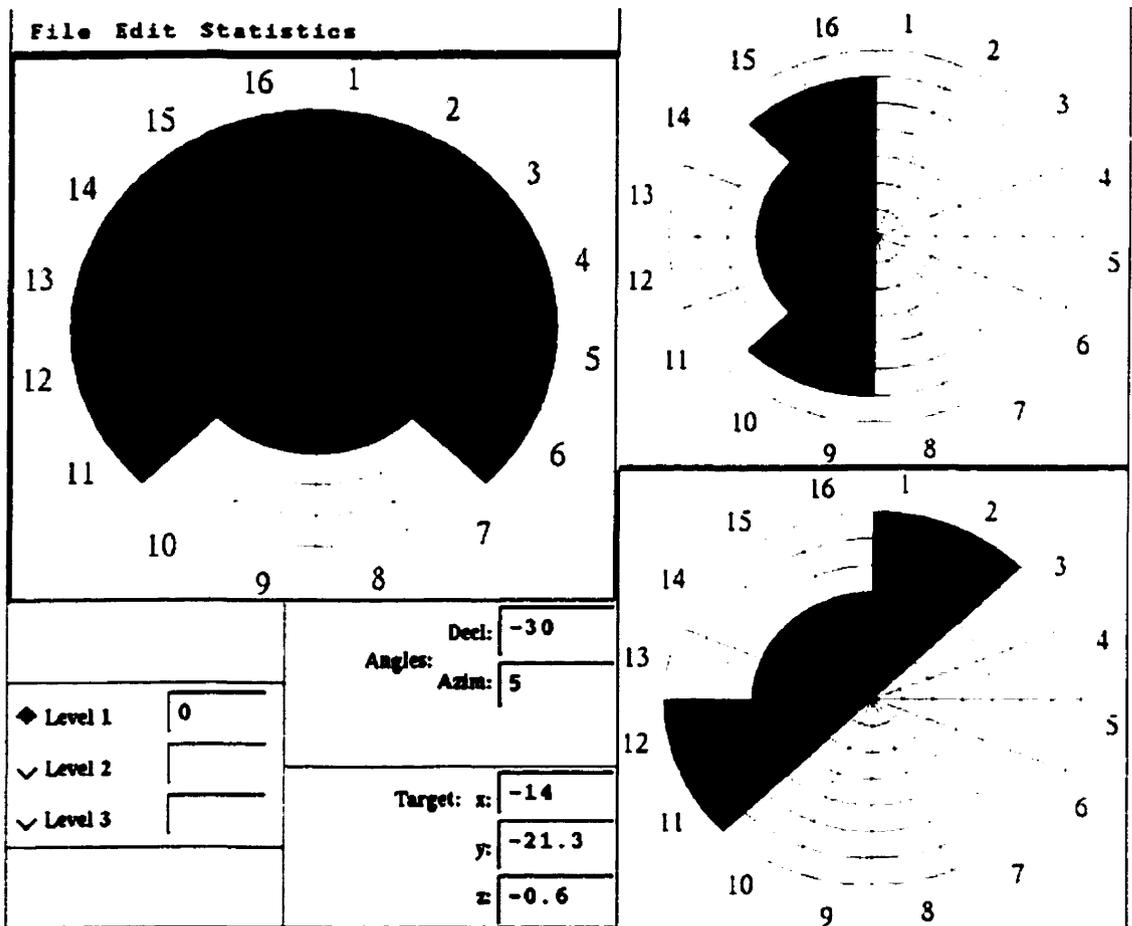


Figure 5.5: The Lesion Modeling graphical interface (left).

On the right, two examples of other possible excision maps.

Chapter 6

Clinical Considerations

We present in the next sections the step-by-step clinical procedure for each case, from the pre-operative scan acquired one week before the operation until inspection of the post-operative scan in the days following the procedure.

6.1 Step-By-Step Clinical Procedure

6.1.1 Pre-operative steps

The pre-operative MRI scan is usually acquired one week before the actual operation takes place. The non-linear transformation that relates the MRI model to this pre- op scan is computed automatically overnight. The atlas is then resampled (or transformed) accordingly to the result of the non-linear fitting algorithm and the “goodness of fit” is inspected visually in regions where the MRI scan possesses sufficient contrast to do so.

6.1.2 After the frame is attached to the patient

Once the frame is put in place, the stereotactic MRI scan as well as the ventriculography are performed. While the neurosurgeon localizes the target point using ventriculograms and the MRI

images, the computer calculates the affine transformation (a simple reorientation) between the stereotactic scan just obtained and the pre-op scan from the week before. Since the atlas is already in correspondence with the pre-op scan, a simple realignment of the atlas is needed according to the computed linear transformation. This step requires about fifteen minutes of CPU time on our machine (SGI R4400 150 MHz). Once the atlas is resampled to fit the stereotactic scan, we perform the semi-automatic segmentation of surfaces representing the different structures that the neurosurgeon wishes to see in the 3D-window. This step takes about half an hour as well. An additional five to ten minutes is needed to perform the MRI space to frame space registration using the fiducial markers appearing on the MRI images.

6.1.3 During the operation

The computer system is then brought to the OR and the VIPER software is activated. Both the stereotactic MRI scan and the co-registered atlas are superimposed in the 2D panels, while the 3D surfaces are visible in the 3D window. The first part of the operation is the stimulation stage (see Chapter 2). During this part, a model of the stimulator is displayed on the screen and correspondence between the position of the electrode and the sensory response of the patient is evaluated. In doing so, the neurosurgeon can verify whether the information found on the computer is reliable or not. From the precise location of critical structures found during the stimulation stage, the neurosurgeon can make a first attempt in deciding on the protocol to use for the leukotome (Chapters 2 and 5). An assistant then creates a modeled lesion that corresponds to this first protocol. Inspection of the relative position of the modeled lesion and of the target volume states whether or not the protocol should be used. A final verification of the position of the leukotome with respect to the atlas in the 3D window in stereoscopic mode is performed before moving to the excision itself. This completes the IGNS work in the OR.

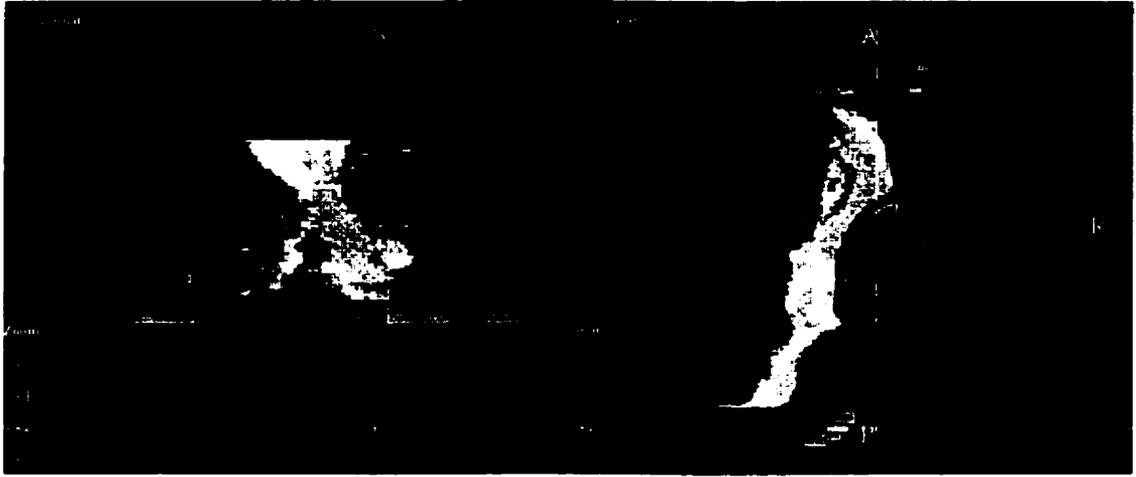


Figure 6.1: 2D views of a post-operative MRI with superimposition of the atlas.

The lesion site can be seen as a dark spot inside the outlined thalamus.

6.1.4 Post-operative steps

In the two days following the operation, a post-operative MRI scan is acquired. At this point, the stereotactic frame has been removed from the patient. On this scan, the site of the lesion is clearly seen. Once the computer algorithm has found the affine transformation (reorientation) that is required to go from the pre-op scan to the post-op scan, we are able to resample the atlas into the post-operative MRI space. The site of the actual lesion can then be compared with the target volume as seen on the atlas (Fig. 6.1).

6.2 Clinical Results

At the time of this writing, eight patients have undergone operations making use of these tools. Six patients required thalamotomies, while only one pallidotomy was indicated. The remaining case was a permanent stimulator placement on a patient who had already undergone two unsuccessful operations in another hospital (one thalamotomy and one pallidotomy). The modeled leukotome and the virtual lesion generator were of no use for this case, but the VIPER platform was found to be useful for the stimulation stage.

At this stage the VIPER platform is used by the surgeon as an investigational tool only, and is always used in conjunction with the standard procedures and tools (ventriculography, intra-operative x-rays, stimulation run...) for these operations.

The goal of our work was not to replace the existing neurosurgical environment, but rather to improve it. While the conventional tools provided sufficient information and precision to allow for this operation to be performed, they sometimes forced the neurosurgeon into very demanding mental work in reconstructing the geometry of structures that cannot normally be seen. In this sense our tools can facilitate the work of the neurosurgeon, providing him with most of the geometrical information he needs. The only task for which the information obtained with the VIPER software could be misleading is in the positioning of the atlas and of the modeled neurosurgical tools with respect to the "real" anatomy, but this is verified through the stimulation procedure.

In some situations, the information provided by the conventional tools is insufficient to discriminate between which of two similar adjustments of the leukotome should be used. This discrimination is never critical to the patient's health, but it is believed that being able to discriminate between the two choices could slightly improve the success of the operation in decreasing the Parkinson's symptoms. In two of the eight cases where the VIPER platform was used, the neurosurgeon was able to discriminate between two a priori equivalent adjustments of the leukotome from the supplementary information that he received from our system.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

The tools that we have developed are strongly clinically-oriented. The goal of our work was to bring computer guidance for thalamotomies and pallidotomies into the operation room. We were fortunate in that a number of conventional IGNS tools simply had to be adapted to our needs (neuro-navigation, multi-modality imaging, registration between MRI space, stereotactic frame space and physical space). Nevertheless much thalamotomy-specific work had to be performed (creation of the atlas, registration of the atlas, leukotome modeling in image space, lesion modeling...). Similar IGNS systems have been developed by other people [31], but to our knowledge none of them included a completely non-linearly deformable atlas, and a lesion modeling algorithm such as ours.

Computer guidance for thalamotomy and pallidotomy was needed essentially to simplify the neurosurgeon's work in understanding the position of the neurosurgical tools with respect to the anatomy. Dr. Sadikot, who performed all eight operations where the VIPER platform and its tools were used, expressed his satisfaction by making the use of our platform part of the routine procedure for thalamotomy and pallidotomy.

7.2 Future Work

Atlas improvement: The interpolation of the acetate contours from the Schaltenbrand atlas (transverse series) to create a volumetric atlas is not perfect. Misalignment of the slices in the Schaltenbrand atlas itself gave rise to a jagged interpolation between certain slices. Even though the manual point-to-point tagging helped in solving part of the problem, some manual editing on the atlas would improve the precision of the atlas description of deep brain structures [31].

Volume Rendering in VIPER: Even though the coding for volume rendering is near completion, its implementation in the VIPER environment is not yet finished. As stated in Chapter 4, volume rendering of the MRI data set is indicated in the task of providing context to the surface rendered objects representing the thalamic nodes.

Introduction of the ventriculograms and intra-operative X-rays into VIPER: The ventriculograms are already scanned into a computer in order to be used by the computerized ventriculography targeting program developed by Roch Comeau in our lab (see Chapter 2). This will be facilitated by the integration of an inter-operative Digital Fluoroscopic unit with VIPER. It would then be useful to implement Roch Comeau's program in VIPER and to be able to compare at any time the position of the stimulator or the leukotome on the X-rays and the modeled version proposed in VIPER.

Creation of a probabilistic functional atlas: Functional data coming from the stimulation runs can be used to create a probabilistic functional atlas of the regions where the stimulation is usually done, i.e. the Vci nucleus, Vce nucleus and internal capsule. For each of the fifty cases where Dr. Sadikot performed either thalamotomy or pallidotomy, the exact setting of the stimulator with respect to the stereotactic frame and the corresponding response of the patient were registered and kept in a file. The MRI scans are archived as well. By transforming the atlas *a posteriori* onto these scans and by identifying the position of the stimuli with respect to the atlas, a probabilistic functional map of responses to stimuli could be superimposed on the

atlas. This additional information could be used by the neurosurgeon to predict the location of the point to stimulate. The probabilistic atlas in itself would present taxonomic interest to the neurology field as well, and add to the functional database being compiled by the ICBM project. Similar work has already been reported [4, 31, 39], but not with a non-linearly deformable atlas such as ours.

Bibliography

- [1] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5):698–700, September 1987.
- [2] G. H. Barnett, D. W. Kormos, C. P. Steiner, et al. Frameless stereotaxy using a sonic digitizing wand: Development and adaption to the picker vista medical imaging system. In R. J. Maciunas, editor, *Interactive Image-guided Neurosurgery*, pages 113–119. American Association of Neurological Surgeons, 1993.
- [3] G. Bertrand. *Computers in stereotaxic surgery*, chapter 19, pages 364–371. Georg Thieme Verlag, 1982.
- [4] G. Bertrand, H. Jasper, and A. Wong. Microelectrode study of the human thalamus. functional organization of the ventro-basal complex. *Confin Neurol*, 29:81–86, 1967.
- [5] G. Bertrand, H. Jasper, A. Wong, and G. Mathews. Microelectrode recording during stereotactic surgery. *Clinical Neurosurgery*, 16:328–355, 1969.
- [6] G. Bertrand, A. Olivier, and C. J. Thompson. Computer display of stereotaxic brain maps and probe tracts. *Acta Neurochirurgica (suppl)*, 21:235–243, 1974.
- [7] 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.
- [8] M. Brain. *Motif programming. The essentials... and more*. Digital Press, 1992.

- [9] D. L. Collins. *3D model-based segmentation of individual brain structures from magnetic resonance imaging data*. PhD thesis, McGill University, Montreal, Canada, 1994.
- [10] D. L. Collins, W. Dai, T. S. Marrett, and A. C. Evans. Three-dimensional non-linear warping of a computerized volume-of-interest brain atlas for morphometric analysis. In *Information Processing in Medical Imaging*, Wye, UK, July 1991. poster session.
- [11] D. Louis Collins, Terry M. Peters, and Alan C. Evans. An automated 3D non-linear image deformation procedure for determination of gross morphometric variability in human brain. In *Proceedings of Third Conference on Visualization in Biomedical Computing*, volume 2359, pages 180–190. SPIE, Oct. 4-7 1994.
- [12] Bruce Davey, Roch Comeau, Chris Gabe, Andre Olivier, and Terence Peters. Interactive stereoscopic image-guided neurosurgery. In *Proceedings of SPIE Medical Imaging*, volume 2164, pages 167–176. SPIE, 1994.
- [13] Bruce Davey, Patrice Munger, Roch Comeau, Laura Pisani, Andre Olivier, and Terry Peters. Applying stereoscopic image visualization to neurosurgical image guidance. In Anthony M. DiGioia, Takeo Kanade, and Russell Taylor, editors, *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, pages 264–271, Pittsburgh, Penn., September 22-24 1994. Shadyside Hospital.
- [14] Bruce L. K. Davey, Roch M. Comeau, Patrice Munger, Laura J. Pisani, Daniel Lacerte, Andre Olivier, and Terence M. Peters. Multimodality interactive stereoscopic image-guided neurosurgery. In Richard A. Robb, editor, *Proceedings of Third Conference on Visualization in Biomedical Computing*, volume 2359, pages 526–536. SPIE, SPIE, Oct. 4-7 1994.
- [15] J. D. Foley, A. Dam, S. K. Feiner, and J. F. Hughes. *Computer graphics, principles and practice*. Addison-Wesley, 1990.
- [16] R. L. Galloway and R. J. Maciunas. An articulated localizing arm for neurosurgical use. In R.J. Maciunas, editor, *Interactive image-guided neurosurgery*, pages 159–168. American Association of Neurological Surgeons, 1993.

- [17] R. L. Galloway Jr., R. J. Maciunas, and C. A. Edwards II. Interactive image-guided neurosurgery. *IEEE Transactions on Biomedical Engineering*, 39(12):1226–1231, Dec 1992.
- [18] R. Hassler. The influence of stimulation and coagulation in the human thalamus on the tremors at rest and its physiopathologic mechanism. In *Proceedings of the second international congress of neuropathology*, London, 1955.
- [19] C. J. Henri, D. L. Collins, and T. M. Peters. Multi-modality image integration for stereotactic surgical planning. *Medical Physics*, 18(2):167–177, Mar/Apr 1991.
- [20] Chris Henri. Application of stereoscopic digital subtraction angiograms to stereotactic neurosurgery planning. Master's thesis, McGill University, Montreal, September-October 1989.
- [21] C.J. Holmes, R. Hoge, D.L. Collins, and A.C. Evans. Posthoc enhancement of mr signal using automatic registration. In *2nd International Conference on Functional Mapping of the Human Brain*, page 240, Boston, June 1996.
- [22] V. Horsley and R.H. Clarke. The structure and functions of the cerebellum examined by a new method. *Brain*, 31:45–124, 1908.
- [23] G.N. Hounsfield, J. Ambrose, B.J. Perry, and C. Bridges. Computerized transverse axial scanning (tomography). *British Journal of Radiology*, 46:1016–1051, 1973.
- [24] A. Kato, T. Yoshimine, T. Hayakawa, Y. Tomita, T. Ikeda, M. Mitomo, K. Harada, and H. Mogami. A frameless, armless navigational system for computer-assisted neurosurgery. technical note. *Journal of Neurosurgery*, 74(5):845–9, May 1991.
- [25] M. Kilgard. *OpenGL Programming for the X Window system*. Addison-Wesley Developers Press, 1996.
- [26] P. C. Lauterbur. Image formation by induced local interactions: Examples using nuclear magnetic resonance. *Nature*, 242:190–191, 1973.

- [27] L. Leksell. A stereotaxic apparatus for intracerebral surgery. *Acta Chir Scand*, 99:229–233, 1949.
- [28] D. McDonald. *Display - Program for display and segmentation of surfaces and volumes*. McConnell Brain Imaging Centre, 1995.
- [29] R. Meyers. The modification of alternating tremors, rigidity and festination by surgery of the basal ganglia. *Res Pub Assoc Res Nervous Mental Disorders*, 21:602–665, 1942.
- [30] P. Neelin, J. Crossman, D. J. Hawkes, Y. Ma, and A. C. Evans. Validation of an MRI/PET landmark registration method using 3D simulated PET images and point simulations. *Computerized Medical Imaging and Graphics*, 17(4-5):351–6, Jul-Oct 1993.
- [31] W. L. Nowinski, T. T. Yeo, and G. L. Yang. Atlas-based system for functional neurosurgery. In Yongmin Kim, editor, *SPIE Vol. 3031 Medical Imaging 1997: Image Display*, Singapore, 1997. Tan Tock Seng Hospital.
- [32] 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 MRI images of the brain. *Journal of Computer Assisted Tomography*, 13(1):20–26, 1989.
- [33] T. M. Peters, B. L. K. Davey, P. Munger, R. M. Comeau, A. Evans, and A. Olivier. Three-dimensional multi-modal image-guidance for neurosurgery. submitted to *IEEE Transactions on Medical Imaging*, 1994, 1994.
- [34] T. M. Peters, C. J. Henri, A. Olivier, and L. Lemieux. Experience with a computerized stereoscopic workstation for neurosurgical planning. In *Proceedings of Conference on Visualization in Biomedical Computing*. IEEE, May 1990.
- [35] G. Schaltenbrand and P. Wahren. *Introduction to stereotaxis with an atlas of the human brain*, volume II. Georg Thieme Verlag, 1977.
- [36] Robin Sibson. Studies in the robustness of multidimensional scaling: Procrustes statistics. *J. R. Statist. Soc.*, 40(2):234–238, 1978.

- [37] E.A. Spiegel, H.T. Wycis, and M. Marks. Stereotaxic apparatus for operations on the human brain. *Science*, 106:349–350, 1947.
- [38] J. Talairach, M. David, P. Tournoux, H. Corredor, and T. Kvasina. *Atlas d'Anatomie Stéréotaxique*. Masson et Cie, Paris, 1957.
- [39] C. J. Thompson, T. L. Hardy, and G. Bertrand. A system for anatomical and functional mapping of the human thalamus. *Comput Biomed Res*, 10:9–24, 1977.
- [40] A. E. Walker. *Classification of movement disorders*, chapter 31, pages 503–509. Georg Thieme Verlag, 1982.
- [41] A. E. Walker. *Classification of movement disorders*, chapter 33, pages 515–521. Georg Thieme Verlag, 1982.
- [42] E. Watanabe. The neuronavigator: A potentiometer-based localizing arm system. In R.J. Maciunas, editor, *Interactive image-guided neurosurgery*, pages 135–147. American Association of Neurological Surgeons, 1993.
- [43] S. Zhai, W. Buxton, and P. Milgram. The partial-occlusion effect: utilizing semitransparency in 3d human-computer interaction. *ACM Transactions on Computer-Human Interaction*. 3(3):254–284, September 1996.