

An Integrated Visualization System for Surgical Planning and Guidance using Image Fusion and Interventional Imaging

David T. Gering¹, Arya Nabavi², Ron Kikinis², W. Eric L. Grimson¹, Noby Hata², Peter Everett², and Ferenc Jolesz² and William M. Wells^{1,2}

¹ MIT AI Laboratory, Cambridge MA, USA
gering@ai.mit.edu

<http://www.ai.mit.edu/people/gering>

² Brigham and Women's Hospital, Harvard Medical School, Boston MA, USA

Abstract. We present a software package which uniquely integrates several facets of image-guided medicine into a single portable, extendable environment. It provides capabilities for automatic registration, semi-automatic segmentation, 3D surface model generation, 3D visualization, and quantitative analysis of various medical scans. We describe its system architecture, wide range of applications, and novel integration with an interventional Magnetic Resonance (MR) scanner to augment intra-operative imaging with pre-operative data. Analysis previously reserved for pre-operative data can now be applied to exploring the anatomical changes as the surgery progresses. Surgical instruments are tracked and used to drive the location of reformatted slices. Real-time scans are visualized as slices in the same 3D view along with the pre-operative slices and surface models. The system has been applied in over 20 neurosurgical cases at Brigham and Women's Hospital, and continues to be routinely used for 1-3 cases per week. ¹

1 Introduction

Image-guided surgery systems facilitate surgical planning and analysis by aligning various data sets with information on morphology (MR, CT, MR angiography), cortical function (fMRI), or metabolic activity (PET, SPECT). Systems for surgical guidance additionally provide a means of establishing a correspondence between the fused, pre-operative data and the patient as positioned on the operating table.

1.1 Limitations of Conventional Image-Guided Surgery

The accuracy provided by frame-based and frameless stereotaxy decreases as surgery progresses. The positions of anatomical structures change as the brain

¹ Ron Kikinis and Ferenc Jolesz received partial support from NIH grants P41 RR13218-01, P01 CA67165-03, and R01 RR11747-01A. Eric Grimson received partial support from NSF grant IIS-9610249.

swells when the skull is opened, cerebrospinal fluid (CSF) egresses, or tumor tissue is removed. Greater dynamics accompany the resection of larger or deeper tumors.

Intra-operative imaging is being developed to address these concerns and also to detect unexpected complications during surgery such as hemorrhage. Interventional imaging is typically performed with X-ray fluoroscopy or ultrasound [2]. More recently, technological developments have led to high-resolution, high-contrast 3D interventional imaging, such as interventional CT [8], and MRI [12]. Compared to the other named interventional imaging modalities, an open MR system provides the advantages of high soft-tissue contrast, a lack of radiation exposure to patient and surgeon, a clear definition of resection boundaries, and continuous access to the operative field [7].

1.2 Augmenting Interventional MRI with the 3D Slicer

In order to amplify the benefits of interventional MRI, we propose augmenting the scanning component with computer software that maximizes the interactivity of an image-guided therapy system through focusing on the following five areas:

Image Resolution — Some anatomical structures are difficult to distinguish on interventional MR images, but are clearer on conventional, diagnostic MRI that benefit from a higher magnetic field and longer imaging times.

Imaging Time — For surgical guidance to be interactive, images must be acquired quickly enough to be utilized without disrupting or slowing down the procedure. Fast imaging techniques are being developed, but in general, faster imaging brings lower image quality.

Multi-modal Fusion — Functional and metabolic data that is acquired pre-operatively could deliver increased benefit if integrated with intra-operative, anatomical information.

Faster Localization — Interventional MR provides the capability of planning approach trajectories by maneuvering a tracked wand and collecting images at the rate of 6-20 seconds per image. While this is valuable, an ideal interactive system needs an update rate of 10 frames per second.

3D Visualization — Interventional images are presently two-dimensional — requiring the surgeon to mentally map the 2D images seen on a computer screen to the 3D operating field.

The 3D Slicer is a software package that addresses the aforementioned areas. Image resolution, imaging time, and localization are improved by performing real-time re-slicing of both pre-operative and intra-operative data sets, and displaying them for simultaneous review. Multi-modal information is incorporated through automatic registration. The 3D Slicer features a computer graphics display that offers the flexibility to see the situation from viewpoints not physically possible. It has the ability to “fly through” virtual data to facilitate the understanding of complex situations, and aid in avoiding damage to healthy tissue.

After testing basic concepts in a clinical trial described in [5], the system was redesigned from scratch. The important developments we report here are the

lessons learned in engineering of the system into an effective and frequently-used surgical tool, and the experiences of using the system in a variety of applications.

2 System Architecture

2.1 Development Platforms

We developed the 3D Slicer on top of the OpenGL graphics library using the Visualization Toolkit (VTK) [13] for processing, and the Tcl/Tk scripting language for the user interface. VTK provides a set of objects written in C++ that can be chained together to form a data-processing pipeline. Pipelined processing maximizes interactivity because the output of each stage is stored in memory, and any update from user-interface controls triggers a change at a minimal depth into the pipeline. We added several classes to VTK by deriving them from existing, documented classes, which results in well-defined inputs and outputs. Appropriately coded classes handle multiple data types and multi-threading. We run the 3D Slicer on PCs running Windows and Sun workstations running Solaris.

2.2 Visualization

Reformatting: We visualize volume data (3D array of voxels) through Multi-Plane Reformatting (MPR). A reformatted image is derived by arbitrarily orienting a plane in 3D space, and assigning values to each 2D pixel of the plane by interpolating the 3D voxels of the volume data intersected by the plane. We reformat up to three slices at once with independent orientations, in real time. Slices can be arbitrarily oblique, or orthogonal and oriented relative to either the coordinate frame of the scanner or the tracked pointing device. Some radiological applications insist on minimum interpolation, and so we provided additional orientation options that generate orthogonal slices relative to the data itself rather than the coordinate frame it has been registered to.

Multiple Volumes on the Same Slice: Each reformatted slice may be generated from any of the available data sets. For example, one slice could be reformatted from preoperative SPGR, a second slice could be reformatted from preoperative T2-weighted MRI, and a third slice could display the images that are being scanned in real time. We extended reformatting to slice through both anatomical and functional volumes of data simultaneously. Each slice is the composite of a background layer, foreground layer, and label layer. The background image is typically gray-scaled anatomical data, and the optional foreground can be colored, functional information. The foreground layer is overlaid on the background layer with an adjustable opacity to form the image that is displayed on the slice. As for the third layer, the output of image segmentation is a label map where each voxel is assigned an integer label according to the segmented structure it belongs to. The boundaries of each of these structures can optionally be drawn in color on the image.

Display Windows: There is an image-processing pipeline for each slice that takes 3D volume data as input, computes a fast reformat using integer math to create 2D data, converts from scalar values to screen colors, overlays the foreground layer on the background layer, and draws cursors and annotation. The output images are displayed in 2D image windows (one for each slice) and texture-mapped onto graphics planes that are correctly positioned in 3D space and rendered in a 3D view. We found that users often need to vary their focus between the 2D images and the 3D rendering. We provided three different methods in attempting to create a convenient interface for all users. First, there is a foveal view in the form of a 2D window that provides a magnified version of the area around the cursor. Second, each 2D image may be zoomed independently. Third, the display mode may be set to show a huge 3D view, a large 3D view and the three 256x256 2D images, or all 4 windows set to size 512x512.

Surface Models: Surface models of key anatomical structures can be visualized in the 3D view along with the reformatted slices. Our surgical colleagues like to see a portion of the skin as a landmark, so we allow for the slice planes to selectively clip away the skin model to reveal other unclipped models beneath, such as a tumor or physiologically critical structures like blood vessels, as well as the respective image planes. In cases where the tumor is near the eloquent cortex, such as centers of speech or motor processing, functional models are included to illustrate their proximity to tumor tissue. These models display the 3D structure of a detected correlation of functional response. Each model is colored differently (and consistently between cases), and rendered with adjustable opacity (Figure 1). We also support applications that require varying a model's color along its surface to visualize a property, such as bladder wall thickness, for example.

Trajectory Assistance: In addition to slices and models, the locator (a tracked wand or surgical instrument) is also rendered as a cylinder in 3D space at its correct location and orientation. The software can also provide means for assisted trajectory planning and guidance. Once a low-contrast target is identified using near real-time imaging, it is sometimes difficult to quickly find it again. This can be simply avoided by marking the probe's position with a bright red sphere on the 3D display, so the location can be easily returned to.

Trajectory assistance is achieved through marking the point of entry with one sphere, and the target tissue with another sphere, and rendering a cylinder between the two points. Thus the surgeon has only to align the rendered locator with the target in a video-game-like exercise to find the planned approach.

2.3 Multi-Modal Registration

The 3D Slicer supports rigid, manual registration by allowing the user to specify which volume to move, and then intuitively translate and rotate that data set by clicking and dragging on the images in the 2D windows. The reference volume can be displayed as the background layer of the slices, and the registered volume

can be displayed translucently as the foreground layer. Furthermore, with one button click, the result of manual registration can be used as an initial pose for automatic registration by maximization of mutual information [16]. This method is of general utility, and other implementations of it have performed well in an NIH-sponsored test [6]. The registration capability is implemented as a separate process that communicates with the 3D Slicer through the use of MRML files as described below.

2.4 Volume Editor

Volumetric data can be semi-automatically segmented using the 3D Slicer’s suite of editing tools. Effects such as thresholding, morphological operations, island-removal (erasing small groupings of similar pixels), measuring the size of islands, cropping, and free-hand drawing of polygons, lines, or points can be applied to the data on a 3D or slice-by-slice basis. We found it very helpful to users to allow the slice-by-slice editing to be administered on either axial, coronal, or sagittal slices merely by clicking on the appropriate slice. Each effect may be applied to either the *original* volume, or a *working* volume. Multiple editing effects may be applied to the working volume, and when finished, the working volume may be merged with a *composite* volume to overlay smaller structures onto larger ones.

For example, we segment skin automatically from an MR scan by applying a threshold at the first trough of the histogram (where the noise and signal lobes overlap) and storing the binary output in the working volume. We then remove islands in the working volume, and finish with an erosion and dilation to smooth the edges, as was performed in the making of Figure 1. This working volume is then copied to the composite volume. Next, segmentation of the brain can be performed in the working volume, and all non-zero voxels of the working volume can overwrite the composite volume to form a combination of skin and brain. A strength of our system is that effects can be visualized by overlaying the working volume translucently on the original volume and explored in the 3D view.

Through a modular architecture that exploits C++ inheritance, other researchers can easily add new segmentation tools to this suite.

2.5 Model Generation

The output of the segmentation process is a set of label maps, where pixels take on values corresponding to tissue type. The bounding surfaces of the label maps are extracted and represented as a collection of triangles using marching cubes. Decimation reduces the number of triangles to a quantity that can be more quickly rendered with little observable loss in detail [14]. For example, a typical brain surface is reduced from approximately 500,000 triangles to 150,000.

We found it helpful to have the 3D Slicer call a separate process to generate the models so that sets of many models could be performed as a batch job in the background while the user continues to use the 3D Slicer for other tasks. Model generation can be run on another machine with more memory and the user is emailed when the job is complete so the models can be viewed in the 3D Slicer.

2.6 Measurements

The 3D Slicer facilitates several types of quantitative measurements that may be made in 3D space. One can click on a particular surface model to measure its surface area or volume. Markers can be positioned on the surface models and slice planes by clicking on the 3D view, and then the distances or angles between markers can be measured.

2.7 Image-Guided Surgery

We integrated the 3D Slicer with the interventional MR system (Signa SP, GE Medical Systems, Milwaukee, WI) to form a computerized surgical assistant. The location of the imaging plane can be specified with an optical tracking system (Image Guided Technologies, Boulder, CO), which we will refer to as the *locator*. Light-emitting diodes affixed to a locator are tracked by 3 CCD cameras that are attached to a rail above the operating field. The spatial relationship of the instrument relative to the scanner is reported as both a position and an orientation with an update rate of approximately 10 Hz.

The Signa SP console and the Signa SP imaging workstation have a network connection with a TCP/IP interface by which we added a visualization workstation (Ultra 30, Sun Microsystems, Mountain View, CA) running the 3D Slicer. The 3D Slicer connects through a socket to a server we created on the SP imaging workstation. Whenever the locator's position or orientation updates, or a new image is scanned, the server sends the new data. The server is concurrent so that when the 3D Slicer opens a connection to it, the server spawns a child process to perform the actual serving, and the parent server goes to sleep. This way, if the server crashes, the 3D Slicer merely reconnects to the sleeping parent, and the surgery can continue uninterrupted.

The visualization workstation contains two Sun Creator3D graphics accelerator cards. One drives the 20 inch display placed in the control area of the surgical suite. The other card outputs only the 3D view and no control panels. Its signal is converted to NTSC TV and displayed on an LCD panel in the scanner gantry.

2.8 Medical Reality Modelling Language

Visualizing medical data involves combining various data sets collected in various geometric locations into a single scene, and exploring the scene interactively. A scene is created from a variety of volume data sets, surface models derived from those volumes, and transformations derived from 3D registrations of both the volumes and models. We have found that the proper coordination of these items is easiest to obtain by the use of a hierarchical modeling paradigm as exemplified by the modeling systems and languages of graphics and CAD/CAM.

Toward this end, we created a novel file format: the Medical Reality Modeling Language (MRML). MRML files are not a copy of the data in another format. Instead, a MRML file describes where the data is stored so the data can remain in its original format and location. Second, a MRML file describes how to position

the data sets relative to each other in 3D space. Third, a MRML file describes how to display the data by specifying parameters for rendering and coloring.

A scene is represented in MRML as a tree-like graph where volumes, models, the coordinate transforms between them, and other items are the nodes in the graph. Each node has attributes for specifying its data. Transforms along each branch are concatenated to form a composite transform applied to the leaf node.

3 Results

To evaluate the performance and utility of the 3D Slicer, we have applied it in a variety of applications including guiding biopsies and craniotomies in the operating room, offering diagnostic visualization and surgical planning in the clinic, and facilitating research into brain shift and volumetric studies in the lab. While extensive details of the applications and their procedures will be reported in medical specialty venues, we include illustrative summaries here.

3.1 Pre-operative Planning

Outside of the operating room, the 3D Slicer is routinely relied upon as a stand-alone tool for analysis and planning. A vitally important component of neurosurgical planning is plotting an approach trajectory that avoids critical structures such as the blood vessels and the motor cortex. Tumors can either infiltrate functional tissue or push it aside, and a tumor that intricately invades eloquent cortex can be considered inoperable for the sake of preserving the quality of life rather than longevity. In the case shown in Figure 1, a functional MRI exam was registered to an anatomical exam. Surface models of the skin and tumor were constructed from the anatomical exam, and models of the motor cortex, and areas for auditory verb generation and visual verb generation were constructed from the fMRI. The same three slices that are visible in the 3D view are displayed as cross-sections at the bottom. The volumetric form of the fMRI results are overlaid in color on the grayscale, anatomical images.

This patient had a tumor in Broca's area on the left side of the brain where 96% of speech is generally processed. The functional MRI exam was administered to ascertain whether the tumor tissue had infiltrated the eloquent cortex. The 3D Slicer's integrated 3D view clearly demonstrated that the speech activity migrated to the right side, proving the operability of this lesion. It is interesting to note that although the 3D Slicer was scheduled for use only in pre-operative planning for this patient, the surgeon came out of the operating room halfway through the surgery to view the 3D Slicer again.

3.2 Surgical Guidance

After several clinical trials during development, we have used the 3D Slicer as a navigational tool in 22 cases between January and May of 1999. This initial feasibility phase proved the 3D Slicer to be a stable and reliable application. As

we apply the 3D Slicer on a routine basis, a clinical evaluation of its influence on surgical decision making and resection control is being conducted.

The 3D Slicer is set-up and tested for about 5 minutes before the patient enters the MR. After the patient is positioned within the interventional MRI, an initial 3D volume is acquired. For cases in which pre-operative data sets have been fused, a second rigid registration using mutual information relates the pre-operative data sets to this new volume in order to align them with the coordinate frame of the interventional MR scanner.

The surgeon points at his intended craniotomy at various angles. As the tracked pointer moves within the surgical field, it is rendered in the 3D view, and the reformatted slice planes follow its position, sweeping through the volumes. The surgeon verifies, and in one case to date, alters, the planned approach by visualizing it on the display relative to all the surface models of critical structures. Reformatted images can be automatically generated along this path with the click of one button. A single 2D view cannot reveal all hazards, but when two slices are reformatted in the plane of the locator, yet perpendicular to each other, multiple hazards in 3D space may be seen simultaneously.

Cerebrospinal fluid leakage as well as tumor removal itself causes brain shift. Subsequent acquisitions at various times during the surgery allow adjustments to intraoperative changes. According to the magnitude of the intraoperative changes these acquisitions are carried through (on average 3-5 times during the case) as demanded by the surgeon. The Surgeon sees the location of the pointer's tip on the display, giving him a direct correlation of his visual impression of the tissue and the MR definition of the tissue. This is especially useful for areas that appear normal to the eye but display abnormal signal intensities in the applied sequences (T1 or T2). The Slicer's capability of displaying intraoperative T1 and T2 imaging as well as preoperative multi-modal information allows the comparative evaluation and the subsequent modification of the surgical tactic. The surgeon can observe the changes in position of anatomy on the intra-operative images, while also benefitting from the higher resolution, higher contrast, functional information, and administered contrast agents of the pre-operative images. Since model deformation has yet to achieve a sufficiently dependable accuracy for surgeons, we may stop using pre-operative models once resection is well underway.

The 3D Slicer's validity for planning and guiding resection is highlighted by the histopathology and location of the clinical cases for which it was used. Ten tumors were low-grade gliomas, for which total resection is correlated with a prolonged survival time and presumably even with total cure. Of these, six were located in the patients' dominant left hemisphere, three in the immediate vicinity of the motor cortex, and two close to speech areas. In these cases, precise margin definition is of utmost importance. Visually differentiating tissue in these tumors can be difficult to impossible, whereas T2-weighted images give a good estimate of the tumor extent. The display of these images allows a more careful control of total tumor resection. Of the remaining seven cases with high-grade (astrocytoma III and glioblastoma multiforme) tumors, four were adjacent to

the motor cortex, thus making updated navigation a valuable tool to define the surgical goal and prevent morbidity.

3.3 Volumetric Analysis and Studies of Dynamics

The 3D Slicer is used outside of surgery for quantitative studies that require making measurements on source images and 3D models simultaneously. Applications to date include modeling the female pelvic floor for studying incontinence [3], measurements of muscle mass, and orthopedic range of motion studies.

The open MR presents unprecedented opportunities for scientifically studying the dynamics that accompany intervention. We are using the 3D Slicer to analyze brain shift [11] in an effort to eventually reach a level of understanding that includes predictive power. Then image-guided surgical techniques currently used in the open MR could be transferred to conventional operating rooms.

4 Discussion

Aside from the ongoing benefits reaped from its applications, the 3D Slicer itself contributes to the field of image-guided surgery in the following four ways.

- 1.) **Highly Integrated Analysis and Visualization** — Image-guided surgery systems strive to achieve one or more of the following:
 - Data analysis
 - Surgical planning
 - Surgical guidance
 - Surgical guidance with intra-operative updates

Several existing software packages, such as ANALYZE [9], MEDx [15], and MNI [10], present a more extensive collection of tools for data analysis and planning, but do not support surgical guidance. Systems have been developed to facilitate surgical guidance tracking [4], but they feature leaner support for analysis and planning, and they are not designed to incorporate intra-operative updates. Interleaving the results of several of the currently available systems can be cumbersome and time-consuming — making it impractical for clinical applications.

Our system is uniquely valuable in its integration of all of these facets into a single environment. Such cohesion is required for seamless integration of pre-operative and intra-operative data. Feedback from clinical users of the 3D Slicer suggest that they most appreciate its ability to help them navigate a complicated assembly of information.

- 2.) **Integration with interventional MR** —

The 3D Slicer is the first software system to augment an open-MR scanner with a full array of pre-operative data. The analysis previously reserved for pre-operative data [9] can now be applied to exploring anatomical changes that occur as surgery progresses.

- 3.) **Medical Reality Modeling Language** —
MRML conveniently describing 3D scenes composed of a heterogeneous mixture of scanned volumes, 3D surface models, and the coordinate transforms that express the geometric relationships between them.
- 4.) **Freely Available, Open-Source Software** —
We have made the 3D Slicer freely available as a tool to clinicians and scientists [1]. It has a modular, easily extendable design for other researchers to expand its value in a rapidly evolving field.

References

1. Brigham & Women's Hospital and Harvard Medical School. *Surgical Planning Lab*. <http://www.splweb.bwh.harvard.edu:8000>.
2. R.D. Bucholtz, D.D. Yah, J. Trobaugh, L.L. McDurmont, C.D. Sturm, C. Baumann, J.M. Henderson, A. Levy, and P. Kessman. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. In *Proceedings First Joint CVRMED/MRCAS*, Grenoble France, March 1997.
3. J. Fielding, H. Dumanli, A. Schreyer, S. Okuda, D. Gering, K. Zou, R. Kikinis, and F. Jolesz. MR based three-dimensional modeling of the normal female pelvic floor with quantification of muscle mass. *American Journal of Roentgenology*, In Press, 1999.
4. W.E.L. Grimson, M. Leventon, G. Ettinger, A. Chabrerie, F. Ozlen, S. Nakajima, H. Atsumi, R. Kikinis, and P. Black. Clinical experience with a high precision image-guided neurosurgery system. Boston, November 1998. Springer-Verlag.
5. N. Hata, T. Dohi, R. Kikinis, F.A. Jolesz, and W.M. Wells III. Computer assisted intra-operative MR-guided therapy: Pre and intra-operative image registration, enhanced three-dimensional display, deformable registration. In *7th Annual meeting of Japan Society of Computer Aided Surgery*, pages 119–120, 1997.
6. J. Fitzpatrick J. West. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21:554–566, 1997.
7. F.A. Jolesz. Image-guided procedures and the operating room of the future. *Radiology*, 204:601–612, May 1997.
8. L.D. Lunsford, R. Parrish, and L. Albright. Intraoperative imaging with a therapeutic computed tomographic scanner. *Neurosurgery*, 15:559–561, 1984.
9. Mayo Clinic. *ANALYZE Software*. http://www.mayo.edu/bir/analyze/ANALYZE_Main.html.
10. MNI. *MNI Software*. <http://www.bic.mni.mcgill.ca/software/>.
11. A. Nabavi, D. Gering, R. Pergolizzi, N. Hata, R. Kikinis, F. Jolesz, and P. Black. Shift happens. Submitted for publication.
12. J.F. Schenk, F.A. Jolesz, P.B. Roemer, et al. Superconducting open-configuration MR imaging system for image-guided therapy. *Radiology*, 195:805–814, 1995.
13. W. Schroeder, K. Martin, and W. Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3-D Graphics*. Prentice Hall, NJ, 1996.
14. W. Schroeder, J. Zarge, and W. Lorensen. Decimation of triangle meshes. *Computer Graphics*, 1992.
15. Sensor. *MEDx Software*. http://www.sensor.com/medx_info/medx_docs.html.
16. W.M. Wells III, P.A. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51, 1996.

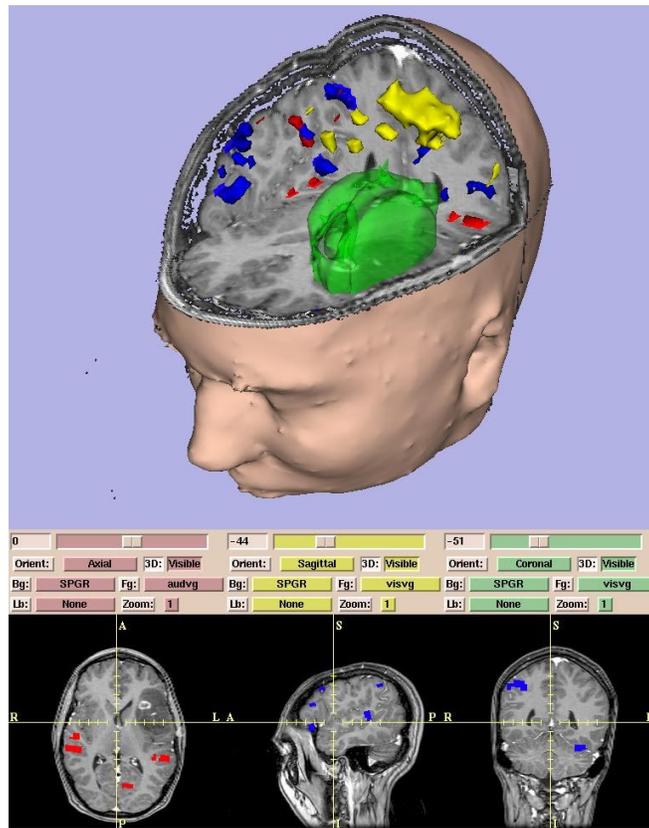


Fig. 1. Surface models of skin, tumor (green), motor cortex (yellow), auditory verb generation (red), and visual verb generation (blue), are integrated with slices.

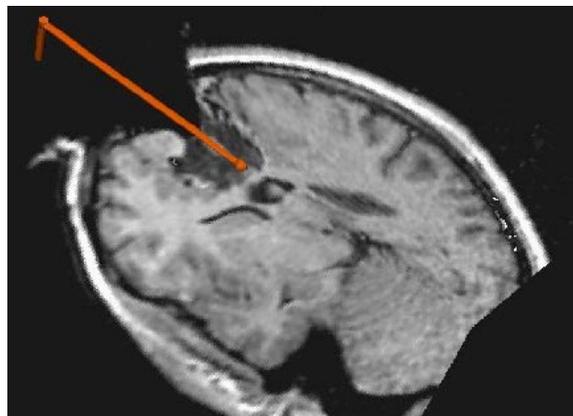


Fig. 2. The surgeon leads a probe through the surgical cavity to check for tumor remnants while observing the 2 orthogonal reformatted images.