The VU-DBS project: integrated and computer-assisted planning, intra-operative placement, and post-operative programming of deep-brain stimulators.

Benoit M. Dawant[a], Pierre-Francois D'Haese[a], Srivatsan Pallavaram[a], Rui Li[a], Hong Yu[b], John Spooner[b], Thomas Davis[c], Chris Kao[b], and Peter E. Konrad[b]

[a] Department of Electrical Engineering and Computer Science, [b] Department of Neurosurgery, [c]Department of Neurology,Vanderbilt University, Nashville, TN 37235, USA;

ABSTRACT

Movement disorders affect over 5,000,000 people in the United States. Contemporary treatment of these diseases involves high-frequency stimulation through deep brain stimulation (DBS). This form of therapy is offered to patients who have begun to see failure with standard medical therapy and also to patients for which medical therapy is poorly effective. A DBS procedure involves the surgical placement, with millimetric accuracy, of an electrode in the proximity of functional areas referred to as targets. Following the surgical procedure, the implant, which is a multi-contact electrode is programmed to alleviate symptoms while minimizing side effects. Surgical placement of the electrode is difficult because targets of interest are poorly visible in current imaging modalities. Consequently, the process of implantation of a DBS electrode is an iterative procedure. An approximate target position is determined pre-operatively from the position of adjacent structures that are visible in MR images. With the patient awake, this position is then adjusted intra-operatively, which is a lengthy process. The post-surgical programming of the stimulator is an equally challenging and time consuming task, with parameter setting combinations exceeding 4000. This paper reports on the status of the Vanderbilt University DBS Project, which involves the development and clinical evaluation of a system designed to facilitate the entire process from the time of planning to the time of programming.

Key words: Computer-aided surgery, registration, deep brain stimulators

1. INTRODUCTION

DBS implantation and programming is a complex task that involves neurosurgeons, electrophysiologists, and neurologists. A typical procedure is sequential. Patients are screened and those who are selected are implanted. The surgical procedure involves a planning phase during which approximate targets are selected because the exact location of these targets cannot be determined with current imaging technology. With the patient awake and using stereotactic techniques, recording and stimulating electrodes are inserted during the surgical procedure. These are used to map regions surrounding the pre-operative target to determine the optimal implant location. When the final positions are selected, the stimulation leads are removed and the DBS leads are inserted to a depth such that the centroid of the four electrodes of each stimulator (Figure 1) is coincident with the final position of the electrode on the respective unipolar stimulation lead. The proximal ends of the leads are then anchored to the skull, and buried beneath the scalp. At our institution, we use a miniature, custom-made, frame —the StarFix microTargeting® Platform (FDA 510(K), Number: K003776, Feb 23, 2001, FHC, INC; Bowdoin, ME, USA)[1] instead of a standard stereotactic frame (see Figure 2) to perform the procedure. As shown in this figure, the frame is

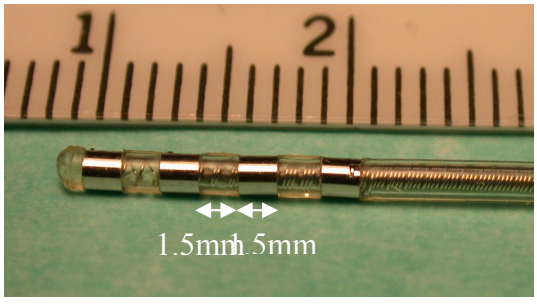

Figure 1: Medtronic #3387 quadripolar lead® (Medtronic, Minneapolis, MN). Each silver band is one electrode. The numbers on the ruler indicate centimeters.

Medical Imaging 2007: Visualization and Image-Guided Procedures, edited by Kevin R. Cleary, Michael I. Miga, Proc. of SPIE Vol. 6509, 650907, (2007) · 1605-7422/07/$18 · doi: 10.1117/12.711149

mounted on bone-implanted anchors. This frame is custom-designed approximately one week before surgery using the location and orientation of the anchors, surgical target coordinates, user-defined entry points, as well the coordinates of the AC (Anterior Commisure), PC (Posterior Commisure), and of one point on the midsagittal plane. The platform is built based on CT coordinates but structures of interest are visible in the MR images, which requires registering the CT and MR images (typically T1-weighted and T2-weighted image volumes). Once a plan is created, a design file is produced and sent to FHC for production. The custom-built platform is sent back to Vanderbilt a few days prior to surgery. The day of surgery, it is mounted on the anchors, and a microdrive is attached to the platform to guide the placement of recording and stimulating electrodes. The StarFix platform has several advantages over the traditional frame: (1) The image acquisition and the target planning can be done prior to the day of the surgery, (2) the smaller and lighter frame allows freedom of movement to the patient, who for movement evaluation is awake during the procedure, and (3) Bilateral implantations can be performed during one procedure. A few days after the procedure, a stimulator is surgically implanted and connected to the permanent DBS lead (Figure 1).

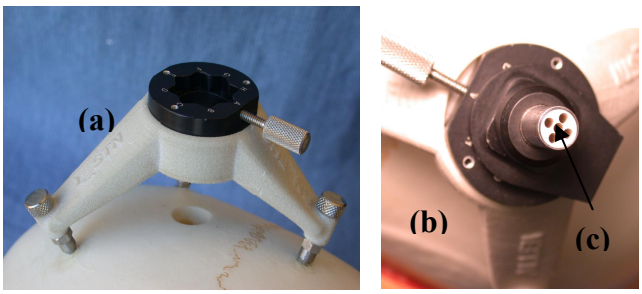

Figure 2: Starfix platform. (a) Mounted on phantom skull. (b) With guide-tubes in place. The central tube guide (c) is aimed at the initially chosen target. The other tubes, each at a distance of 2mm from the central tube are used for lateral adjustments. The cluster of 5 tubes can itself be shifted radially by 3 mm (not shown).

About a month after the procedure, the stimulator is programmed by a neurologist. The goal of DBS programming is to find the combination of voltage, pulse width, frequency, and active contact configuration that generates the greatest symptomatic relief while producing the least adverse effects. With the patient off medication, the neurologist evaluates each electrode contact sequentially in a monopolar configuration to determine the contact producing the best clinical response. Frequency and pulse width are typically kept at constant settings of 130-180 Hz and 60-120 μs respectively. Amplitude is steadily increased until first, benefit of symptoms occurs, and then, adverse effects occur. Repeated neurological examination is performed to assess the efficacy of stimulation. Several minutes are allowed to pass between trials of each contact to allow the effects from the previous stimulation trial to disappear. If a satisfactory result cannot be achieved with monopolar stimulation, more complex configurations such as bipolar, tripolar, or multiple cathodes may be tested. Because of the large number of parameter permutations that could be tried for each DBS patient, the programming process can become very taxing. Furthermore, due to delayed clinical effects, finding the optimal settings may take several trials over many months, which can be frustrating for the patient and clinician.

In a typical clinical setting, a wealth of both population and patient-specific information that could facilitate the process at every stage is not used. This includes statistical atlases that contain population information about optimal target localization, atlases of patients' response to stimulation, or specific information acquired during the surgical procedure that is not available to the neurologist at the time of programming. To address this issue we are developing a system that will transform the overall procedure from a sequential process to a process in which information acquired and used at every step will be available at every other step. This system is composed of several modules: a planning module, an intra-operative module, and a programming module. These modules are linked and communicate with each other through a central Oracle database. This central database can be accessed by the various users through web interfaces that support the tasks to be performed.

The remainder of the paper presents the overall architecture of the system we have developed and implemented, discusses some of the modules in more details, and reports on its use.

2. METHOD

Figure 3 illustrates the various tasks that need to be performed. The process starts as soon as a patient has been scanned. A protocol has been put in place, which permits the automatic transfer of the images from the CT and MR scanners directly to a DICOM client in the School of Engineering Medical Image Processing Laboratory (the MIP laboratory). Once images have been transferred, a web-based interface permits physicians to request the automatic computation of the information required for planning the procedure.

2.1 Creation of planning files

Traditionally, generating a plan for platform creation requires loading the images in a planning software (at our institution the software that is used is Voxim, (IVS Solutions AG, Chemnitz, Germany), registering the various imaging modalities that have been acquired for this patient, and then manually selecting the targets, entry, AC, PC, and mid-sagittal points. Because target points are not easily discernible in current MR images, their selection is typically done based on their AC-PC coordinates, which are adjusted to take into account morphological differences between patients (e.g., large or small ventricles or noticeable asymmetry).

We have proposed an alternative to this approach, which consists in using statistical atlases and non-rigid registration to predict automatically points required for planning. Briefly (more details on the procedure we use can be found in [2,3], we have created MR atlases that contain these points. When a new patient is processed, these atlases are registered to the MR images of the patient and the points of interest are projected from the atlases to the patient volume. We currently use three MR atlases and we have developed strategies that allow us to combine the predictions generated by each of these (see D'Haese *et al.*[3]) .

We have created two types of target point atlases: the first is based on the intra-operative coordinates of the implant (i.e., the position, which appeared to be the optimum during the surgery); the second is based on the most efficient contact, as determined at the time of programming. In both cases, we acquire the coordinates of the target in a number of subjects, register all these subjects to the atlases, project individual target points onto the atlases, and compute the centroid of the projected points. We used these centroids as the ideal target points in each atlas. To localize the AC, PC, and midsagittal points in a patient image, we follow the same procedure, except that these points have been manually localized in the atlases. A study that compares automatic AC and PC predictions to manual selection can be found in Pallavaram *et al*[5].

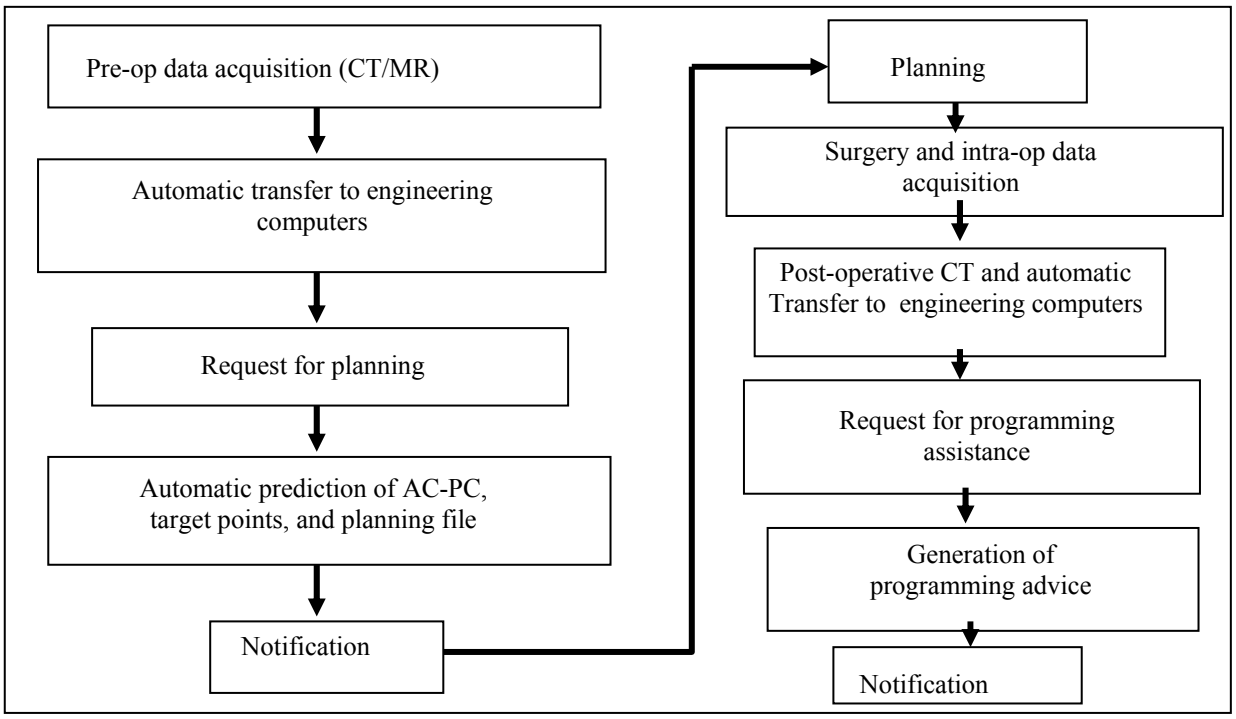

Figure 3: Steps included in the computer assisted DBS procedure at our institution

We have automated the prediction process. Once a request for planning has been logged, the system starts the procedure. This involves (1) registering the CT and MR images of the patient using an MI-based rigid registration algorithm, (2) register non-linearly each of the MR atlases to the patient MR volume using a non-rigid registration algorithm we have developed[6], (3) project the points of interest from the atlases to the patient using the combined transformations, (4) combine the predictions[3], (5) generate an output file that contains the rigid-body transformations and the points of interest, (6) generate an e-mail to the surgeon who requested the plan, and (7) post the information on a secure web-site. The overall process used to take several hours. Parallelization and optimization of the registration algorithms has reduced this time to less than one hour.

2.2 Planning

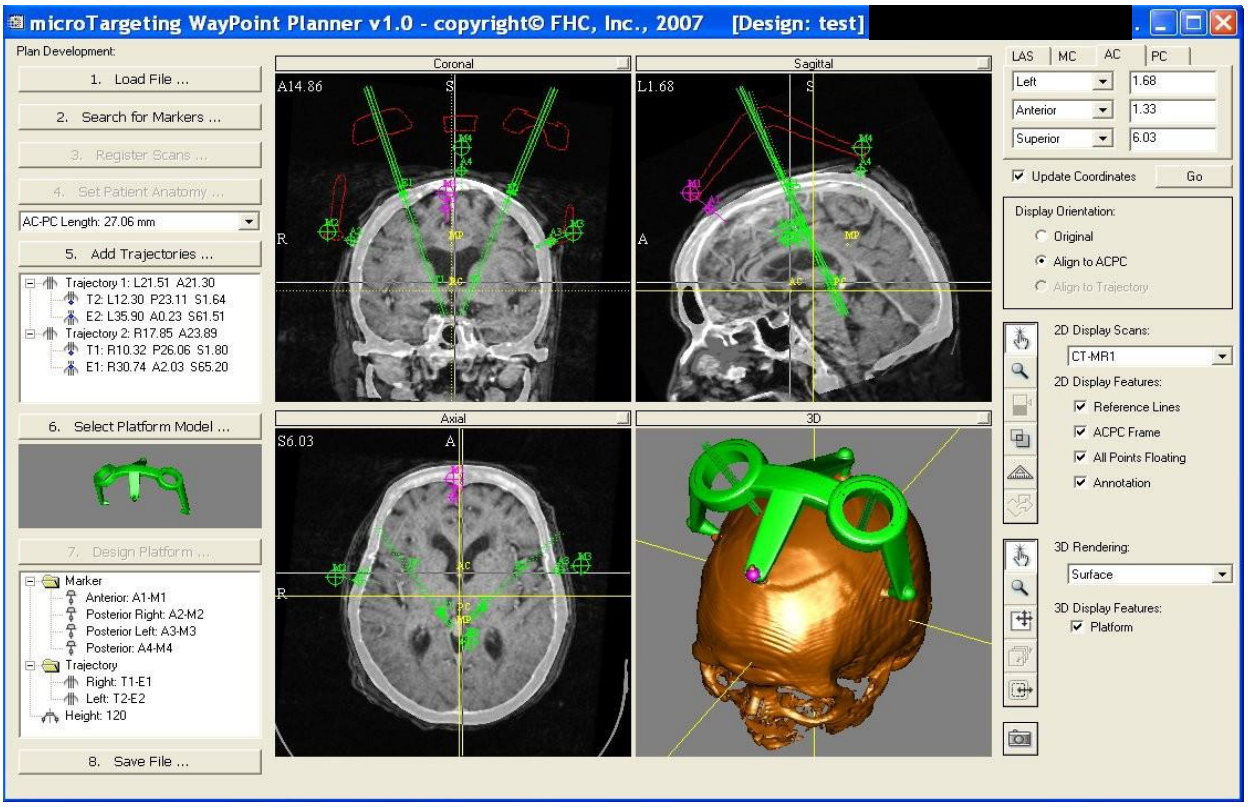

Figure 4: Snapshot of the interface of the system we have developed for planning the procedure

We have collaborated with the company that sells and manufactures the frames (FHC, Inc.) and we have designed and implemented our own planning software to generate the design file. This software permits loading the various DICOM images, registering them to the CT images, automatically localizing the anchors and markers using an algorithm described in Liu et al.[7], manually selecting the AC, PC, midsagittal, entry, and target points, and visualizing the platforms. This software is now in validation use by FHC to create frame design files at our institution. This planning software is also fully integrated with our processing pipeline. The output file generated in the previous step can be read by the planner. Pre-computed transformations can be used to register the various image volumes upon loading, and predicted AC, PC, midsagittal, and surgical target points can be displayed. Planning only requires localizing the anchors and/or markers using the automatic algorithm, verifying and, if needed, adjusting

predicted points, and selecting the appropriate entry point. Figure 5 shows a snapshot of the interface of the mT Waypoint Planner with a bilateral platform.

2.3 Surgery and Intra-operative Data Acquisition

As discussed above, the procedure is iterative. Intra-operatively, micro-electrode recording and stimulating leads are advanced into the patient to the initial target positions determined pre-operatively through the guides visible in Figure 2 . Resting firing frequencies and patient response to stimulations are noted and the target positions are revised. During the procedure, raw MER signals and responses to stimulations are recorded. When possible, raw MER signals are labeled as originating from specific structures such as the STN, Vim, Substantia Nigra, etc. Stimulation-related information includes voltage or current at which the best efficacy was observed, pulse width and frequency, % of efficacy (loss of rigidity or loss of tremor), smallest voltage or current at which side effects were observed, side effect observed, and the body location at which the side effect was observed. The depth position for every measurement, either MER recording or stimulation, is also recorded. After the procedure, intra-operative files are automatically parsed, depth information transformed into x, y, and z coordinates, and the information is stored in the central database as is the final position of the electrode.

2.4 Post-operative CT

As part of the standard procedure, CT scans are acquired post-operatively to verify the position of the DBS leads. These CT scans are also transferred to the Medical Image Processing laboratory and imported in the database. The CT scans are processed and the position of the leads in these post-operative images is determined. This permits the computation of what we refer to as the Electrode Placement Error (EPE), i.e., the difference between the position of the electrode provided by the intra-operative guidance system and the actual position of the electrode. This error may be caused by a number of factors ranging from platform inaccuracy, to bending of the lead, to errors made when substituting the stimulating electrode with the final DBS lead, to brain shift. We have developed an algorithm to detect the location of the leads in the post-op CT, which makes used of intra-operative information. The algorithm starts by registering the pre- and post-operative CT scans. The final implant position (i.e., the position at which the implant is placed intra-operatively) is retrieved from the database. This information is used to place a bounding box in the post-operative CT image volume. The dimensions of this bounding box are chosen in such a way that it covers and extends beyond the four contacts of the electrode. First, a threshold is applied to the region inside the bounding box. Next, a search starts from the most inferior slice in the CT volume that intersects the bounding box until pixels with an intensity value above the threshold are detected (see green line in Figure 5). The centroid of these points is considered to be the bottom corner of the lead. Next, the algorithm jumps 2 slices, ie., typically ~1mm (see red line in the figure) to reach a slice that contains the entire cross-section. The algorithm then progresses in the superior direction for ~15 slices, depending on the implant type. In each slice the centroid of the bright pixels is computed. A vector is computed with these points, which provides the direction of the lead (angle alpha in the figure). Geometric properties for the lead placed in a particular patient are retrieved from the database and used to compute the position of the lead's bottom point. This is done using the trigonometric relations shown in the figure. Once the bottom of the lead has been localized, the point between the two central

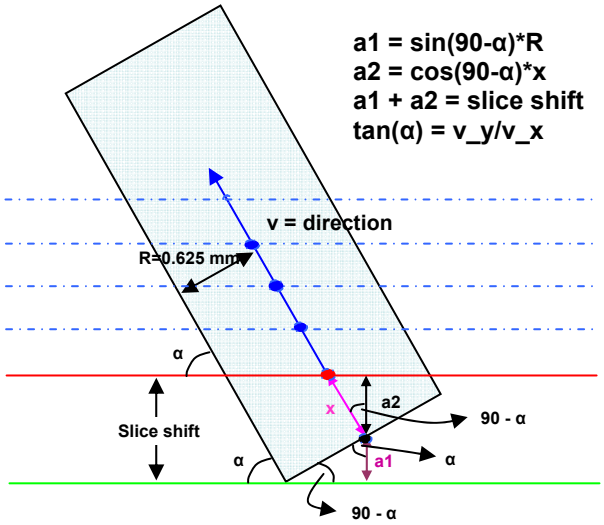

Figure 5: Schematics of the algorithm used to find the position of the DBS lead in post-operative CTs.

contacts, which is compared to the final intra-operative position, is determined. Figure 6 shows results we have obtained with this approach. Preliminary results we have obtained with this method show EPEs up to 3mm.

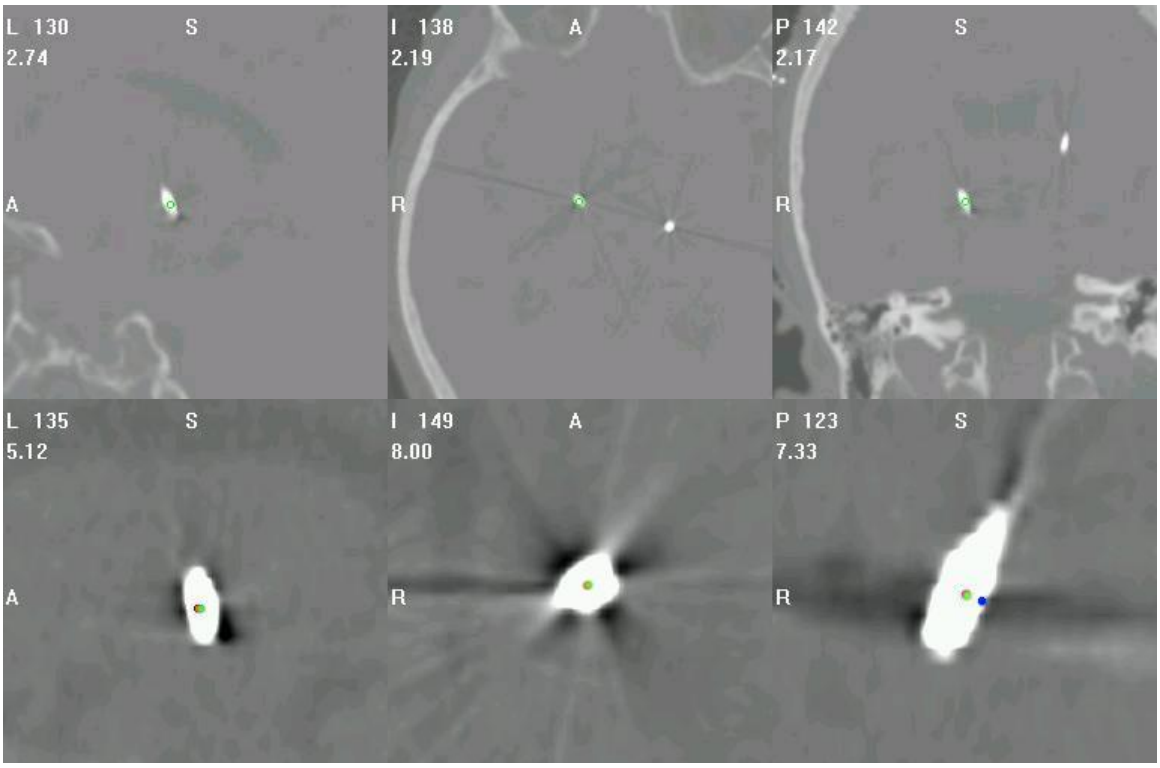

Figure 6: Implant in a CT image. Top panel, sagital, transverse, and coronal views in a CT volume. Bottom panel, same as top but zoomed on the implant. The floating point number shown is the zoom factor. The blue circle in the bottom panels shows the location of the center of the lead determined intra-operatively. The red and green circles show the center of the lead determined the post-operative CT images manually and automatically, respectively

2.5 Generation of Programming Predictions

As discussed above, stimulation testing is performed intra-operatively at multiple points along the planned trajectory to determine the optimal implant position. At each position, stimulation is typically tested with incremental voltages ranging from 0.5 V to 5 V (frequency and pulse width are held constant at 150 μsec and 150 Hz). The effect of the stimulation on the neurological exam is assessed by a neurologist for every voltage increment at every position. The optimal voltage is determined at each position and the loss of rigidity or tremor expressed in percent is recorded for this voltage. We have used the value of three variables measured intra-operatively (loss of rigidity or tremor in %, the voltage required to achieve the desired effect (V), and the therapeutic window (TW), which is the difference between the voltage required to suppress the symptoms and the voltage required to produce side effects) to create a tool that could provide advice to the neurologists at the time of programming. The procedure we have used is as follows. First we register the patient image volume to the atlas. Next, we determine the position of the contacts in the patient volume using either post-op CT or information provided by the intra-operative guidance system. We then define a region of interest for each contact (a cylinder with a radius of 0.675 mm and a length of 1.5 mm). The stimulation data contained in the atlas within this volume of interest is retrieved and used to compute an index that indicates the likelihood that one contact is better than the other. To compute this index, we use a series of heuristics. First, we assume that every point with an efficacy of at least 75% (rated by any of the neurologists) could potentially provide clinically significant symptomatic relief, while points with less than 75 percent are discarded. The likelihood

is then defined as shown in equation (1), in which N is the number of points in the atlas for which we have a measurement in the contact volume. This equation states that we first compute the average therapeutic window over the contact volume. If it is below 2 volts, this contact is considered to be sub-optimal because, practically, 2 volts are needed to provide enough flexibility to fine tune the programming parameters. A contact is better than another if measured efficacy for the points within its contact volume is larger and if the voltage required to suppress the symptoms is smaller.

$$L_j = \frac{1}{N}\left(\sum_{\substack{over\ the \\ contact\ volume}} \frac{Efficacy_i(x,y,z)}{V_i^2(x,y,z)} * \delta(x,y,z) \right) \quad \text{with} \quad \delta(x,y,z) = \begin{cases} 1 & \text{if } avg(TW(x,y,z)) \geq 2 \text{ volts} \\ 0 & \text{otherwise} \end{cases} \qquad (1)$$

Predicted ranking	# of neurologist-selected contacts	%
Rank1	35	57.5
Rank2	13	21.3
Rank3	10	16.4
Rank4	3	4.9
Total	61	100

Table 1: Comparison between contact suggested by the system and contact selected by the neurologist

This index is computed for each contact, and the contacts are ranked. A preliminary retrospectively study performed on 33 STN patients and 61 implants shows the feasibility of our approach. The results of this study are summarized in Table 1. In this table, we correlate the rank of the contact the system predicted to be the best with the contact selected by the neurologist. This table shows, for instance, that in 57.5% of the cases the contact the system predicted to be the best was also chosen by the neurologist.

An alternative use of the intra-operative information we have acquired is the creation of what we call efficacy maps[4]. To create these maps we associate a Gaussian with each point in the atlas for which we have intra-operative data as follows:

$$G(x,y,z) = Efficacy * TW * \frac{1}{V} * \exp\{-\left(\frac{x^2+y^2+z^2}{V^2} \right)\} \qquad (2)$$

A point with a small stimulation voltage and a high efficacy will thus be associated with a curve with a small standard deviation and large amplitude. We then average all the curves to create a probability map. In this average, a point with a large standard deviation has a small impact that extends over a larger region; a point with high amplitude and a small standard deviation has a large local influence. The map can then be projected onto a patient's volume and the lead in this volume superimposed to the map. Figure 7 shows representative results we obtain with this approach. In this figure, white corresponds to a high value in the map and blue to a low value. It suggests that the best contact is the one labeled "therapeutic contact".

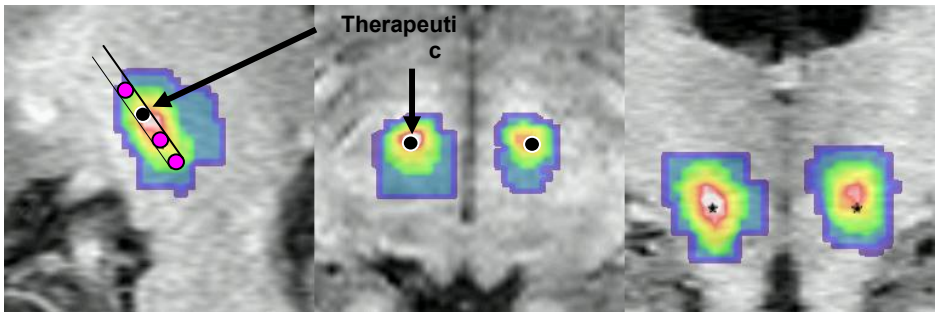

Figure 7: DBS lead superimposed on efficacy map.

2.6 System architecture

Figure 8 shows how the various components of the system interact. One machine hosts the secure web server through which the system is accessed. Another machine hosts the secure Oracle database, which is synchronized

with its clone on a secondary machine for backup purposes. Processing is currently done on a multi-processor machine (2 dual-core Xeon(TM) 3.6 GHz processors, 2 GB of RAM) but the system has been designed to distribute processing tasks across several nodes.

Database users are grouped in teams, with one team per neurosurgeon. Access to the database is accomplished through what is called a virtual private database. Thanks to this mechanism, each team only has access to its own patients. This allows us to manage several groups of physicians while maintaining patient confidentiality. It also lets us customize atlases to reflect the practice of the different groups.

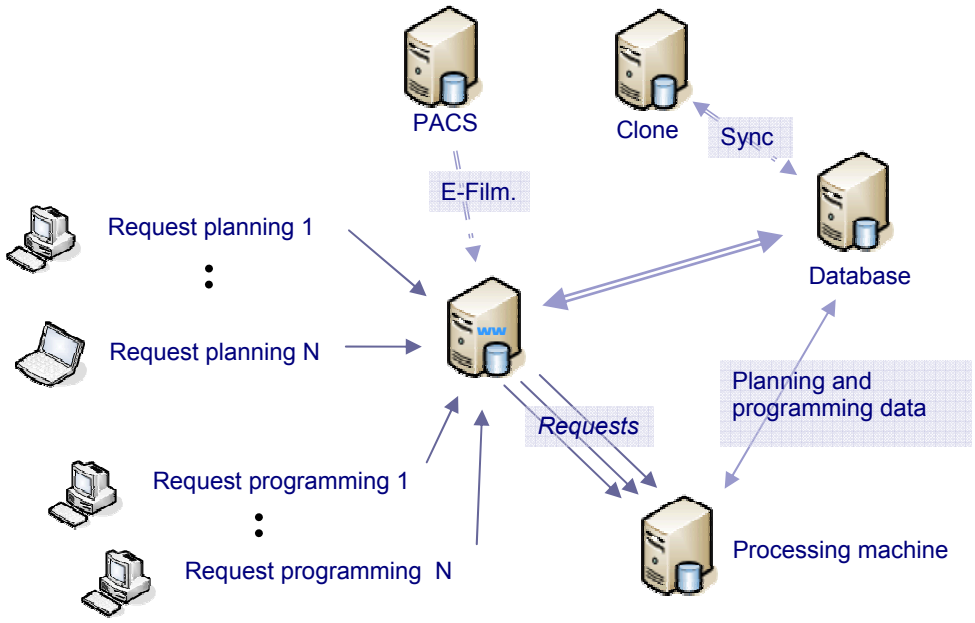

Figure 8: System's architecture

In practice, imaging data needs to be transferred to our system first. This can be done either directly to our DICOM client if a link has been established between this client and the hospital PACS system. If this is not the case, images can be transferred via our web site through a secure applet. Once this is done a user requests a plan for a patient by specifying a patient ID, the desired targets (e.g., Left STN, Right Vim, etc.), the planning preference (i.e., use an atlas made with final implant locations or most efficacious contacts), and the procedure by which (s)he wants to be notified (e.g., e-mail) when predictions are ready.

When a request is received, one member of the engineering team verifies the image quality and also that the images have been entered correctly in the database. When this is done, the planning request is labeled as ready and it is added to a queue on the processing machine. Jobs on the queue are processed sequentially. When a request is processed, the processing machine queries the database for all relevant information (images, voxel size, etc.) and it starts the various registration algorithms required to produce the planning file. Our current implementation is multi-threaded and it can process several atlases at once. Once images are received, prediction files can be produced in less than one hour. When the plans are ready, the requestor is notified that the plan is accessible through the web site.

Programming predictions are handled in a similar way. Requests for programming predictions are made through a secure web-server. When these are ready the requestor is informed and can retrieve the information from the web site.

3. RESULTS

The system has reached the state of a clinical prototype. Predictions are now made for every DBS case performed at our institution. Figure 10 shows the number of predictions that have been made over the years. Prior to 2005, generating predictions required a substantial amount of interactive work

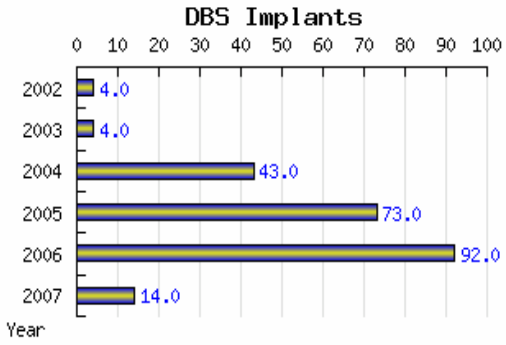

Figure 10: number of predictions performed by our system

and these could not be performed on a regular basis; the processing pipeline we have generated has largely automated the procedure and allows us to support the clinical load. In the summer of 2006, a second neurosurgeon joined the team and we expect a substantial increase in the number of cases. The architecture we have developed is scaleable and should allow us to continue making predictions for every case. As discussed earlier, the structure of the database we have developed permits separating the cases performed by different neurosurgeons, which will permit the creation of individualized atlases. Having two neurosurgeons at our institution permits us to test this feature of the system.

Entering intra-operative information in the database, which involves parsing excel files created during the procedure as well as files retrieved from the MER recording device requires manual intervention. Future versions of the planning system will include an intra-operative component that will connect directly to the database and permit entering this information directly. The component of the system that permits computing the electrode placement error is implemented and works with the database to retrieve the information it needs. It is not yet part of the processing pipeline and it requires manual intervention.

The programming module has also been implemented and works with the database but it is not part of the processing pipeline. Requested predictions need to be processed on a case-by-case basis, which does not yet permit routine, clinical use of this module. A study aimed at determining the best way to use atlases of intra-operative stimulation response is ongoing. This study also analyzes the impact of EPE on the system's prediction. Anecdotal evidence shows that programming predictions obtained with and without taking the EPE into consideration can differ by one contact.

4. DISCUSSION AND CONCLUSIONS

A combination of state-of-the-art registration algorithms, information technology, and a tight interaction between engineers and several clinical departments has permitted the development, implementation, and clinical utilization of what we believe is the first information repository that will permit (1) the development of algorithms, tools, and user interfaces that will facilitate each step of a DBS implantation procedure, (2) easy and complete communications between the various players involved in the process, and (3) constant and immediate feedback between these players (e.g., update surgical targets based on final electrode selection). Because the system is integrated into the clinical flow, the engineering team receives constant feedback from the medical team. For each case we perform, we store the coordinates of the target points predicted by the system, the coordinates selected by the physician when planning the case (this information is produced by the planner), and the final implant position. This allows us to gather statistics on the accuracy of our system and monitor its performance. Based on this information, registration algorithms, which are at the core of the planning procedure can and are still being tuned.

The work we have done so far indicates that retrospective validation of programming predictions, although informative, does not permit drawing hard conclusions. When the system predicts a contact that is different than the one used by the neurologist, it is difficult to assess which one would produce the best clinical results. Integrating the programming module into the pipeline and making it accessible to the neurologists will facilitate the conduct of prospective studies in which both approaches can be compared based on outcome. Inclusion of outcome information in the database is also ng. Currently, our targeting atlases are generated with all the data points we have at our disposal. Whether or not using only patients for which the procedure has been successful would make any difference is under investigation.

In D'Haese *et al.*[2] we present the prototype of an intra-operative user interface. Further development of this interface and its integration with the database will permit to both enter intra-operative information into the database and project population-based information onto the current image volume. This will permit developing algorithms and techniques that use information in the database to provide intra-operative guidance.

The data repository we have developed will also serve as a research tool for a number of scientific questions. For instance, the accuracy and robustness of various rigid and non-rigid registration algorithms can easily be compared. All that is required is to substitute one registration algorithm for the other in the registration pipeline. A number of metrics such as prediction accuracy, speed, or tightness of statistical maps could be used to compare these

algorithms. Different methods by which response to stimulations could be used for programming assistance can also be implemented and compared. The same is true for electrophysiological maps created from MER recordings[2].

We are now looking at opening our planning system to outside groups and we have already tested it with one. Although our system has been tested with the StarFix platform, it can also be used by neurosurgical groups that use a standard stereotactic frame. The key difference between standard frames and the StarFix platform is that, with traditional frames, a scan is acquired the day of surgery with the frame attached; planning is thus done on that same day. This, however, is not a major obstacle. The only thing our system needs is a clinical MR acquired without a frame some time before the surgery (the CT volume is not needed for prediction purposes); these scans are usually part of the standard diagnostic process. With this image volume, predictions can be made ahead of the surgery. The day of the surgery a simple rigid-body registration between the MR volumes with and without the frame would permit projecting all our predictions on the volume acquired the day of surgery.

5. ACKNOWLEDGEMENTS

This work was supported, in part, by grant 1R01EB006136-01 "Computer-aided neurosurgery".

6. REFERENCES

[1] Franck J, Konrad PE, Franklin R, Haer F, Hawksley D. "Starfix : A Novel Approach to Frameless Stereotactic Neurosurgery Utilizing a Miniaturized Customized Pretargeted Cranial Platform Fixture – Technical Description, Unique Features, and Case Reports", Movement Disorders Society, 7th Intl. Congress of Parkinsons Disease & Movement Disorder, Miami, FL, November 2002.
[2] D'Haese PF, Cetinkaya E, Konrad PE, Kao C, and Dawant BM, "Computer-aided placement of deep brain stimulators: From planning to intraoperative", IEEE TMI, 24 (11): 1469-1478, 2005.
[3] D'Haese PF., Pallavaram S., Yu H., Spooner J., Konrad PE., and Dawant BM, "Deformable physiological atlas-based programming of deep brain stimulators: a feasibility study", Proceedings of WBIR 2006, pp. 144-150.
[4] D'Haese PF., Pallavaram S., Niermann K., Spooner J. , Kao Ch. , Konrad PE., and Dawant BM "Automatic Selection of DBS Target Points Using Multiple Electrophysiological Atlases", Proceedings of MICCAI 2005, pp. 427-434, 2005.
[5] Pallavaram S., Yu H., D'Haese P-F, Spooner J, Koyama T, Bodenheimer B, Kao Ch, Konrad PE, and Dawant BM, "Automated selection of anterior and posterior commissures based on a deformable atlas and its evaluation based on manual selections by neurosurgeons", Proc. SPIE Medical Imaging 2007, San Diego, CA, (this issue).
[6] Rohde GK, Aldroubi A, and Dawant BM, "The adaptive bases algorithm for intensity-based nonrigid image registration," *IEEE Transactions on Medical Imaging,* 22(11), pp. 1470-1479, 2003.
[7] Liu X., Cevikalp H., and Fitzpatrick J.M., "Marker orientation in fiducial registration", Proc. SPIE Medical Imaging 2003, San Diego, CA, vol. 5032, 1176–1185 (Feb 2003).