Stereoscopic Display Technologies, Interaction Paradigms, and Rendering Approaches for Neurosurgical Visualization

Jeremy R. Cooperstock and Guangyu Wang

Centre for Intelligent Machines
McGill University
3480 University Street, Montreal
Quebec H3A 2A7, Canada

ABSTRACT

We conducted a comparative study of different stereoscopic display modalities (head-mounted display, polarized projection, and multiview lenticular display) to evaluate their efficacy in supporting manipulation and understanding of 3D content, specifically, in the context of neurosurgical visualization. Our study was intended to quantify the differences in resulting task performance between these choices of display technology. The experimental configuration involved a segmented brain vasculature and a simulated tumor. Subjects were asked to manipulate the vasculature and a pen-like virtual probe in order to define a vessel-free path from cortical surface to the targeted tumor. Because of the anatomical complexity, defining such a path can be a challenging task.

To evaluate the system, we quantified performance differences under three different stereoscopic viewing conditions. Our results indicate that, on average, participants achieved best performance using polarized projection, and worst with the multiview lenticular display. These quantitative measurements were further reinforced by the subjects' responses to our post-test questionnaire regarding personal preferences.

Keywords: Stereoscopy, 3D interaction, medical visualization, user studies

1. INTRODUCTION

The research described in this paper constitutes part of a larger effort to develop a high-performance system for medical data visualization, which supports 3D display of the content and allows similar interaction techniques as would be appropriate for manipulating physical objects. These objectives suggest several requirements, including direct volume rendering for efficient display, integration of multiple data modalities, object tracking, stereoscopic display, and development of the gestural vocabularies as necessary for each task. At each stage of the research, we are conducting experiments intended to identify the salient issues related to each design decision and help make the best choice that is optimized for the needs of the intended users.

Despite the wealth of literature related to stereopsis and interaction paradigms, as reviewed below, there remains the need for a systematic evaluation of interface and display options for interaction with, rather than observation of, medical data. This is a central issue of concern for our application domain of interest, namely neurosurgical visualization. Specifically, we are interested in understanding what display technique is best suited for such a task. Therefore, the first stage of our research deals with the evaluation of different display modalities and input devices for a typical medical visualization task, namely, positioning a biopsy probe along a vessel-free path, as shown in Figure 1. Although we also describe a number of input device options being investigated, for the purpose of this paper, we focus our discussion on the choice between monoscopic displays and three different categories of stereoscopic display, specifically, through our experiment to quantify user performance on the given task.

Further author information: (Send correspondence to Jeremy R. Cooperstock)
Jeremy R. Cooperstock.: E-mail: jer@cim.mcgill.edu
Guangyu Wang: E-mail: gywang@cim.mcgill.ca

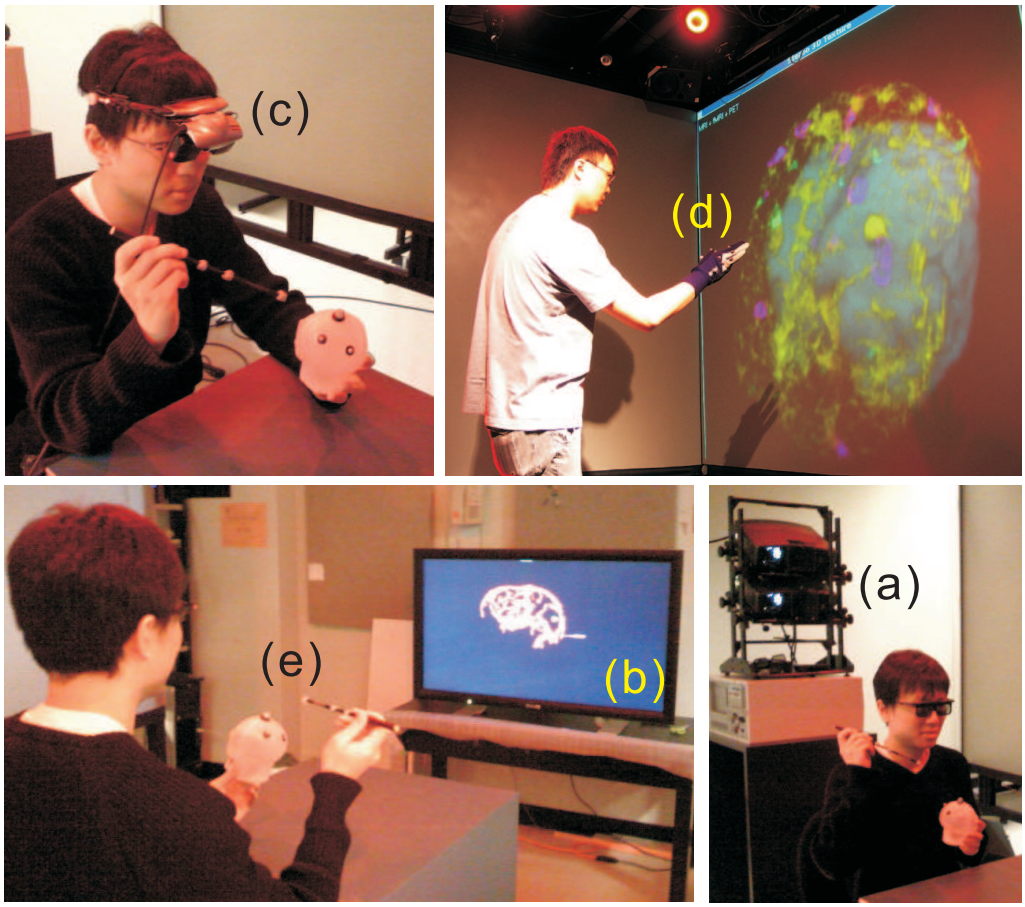

Figure 1. In our system, we provide three display choices: (a) polarized projection, (b) multiview lenticular display, and (c) head-mounted display. Such displays assist users to achieve 3D experience. There are two interaction paradigms: (d) hand gesture and (e) tangible user interface.

2. BACKGROUND

2.1 Display Technologies

Although some surgical training and planning systems provide true 3D visualization capabilities,[1,2] most systems are based predominantly on conventional 2D display. One reason for this may be the cumbersome nature of the technology and the impediments it poses to interaction with other tools and communication with other people.

Stereoscopic viewing can be supported by a number of technologies, including head-mounted display (HMD), two-color anaglyph, polarized stereo projection, LCD shutter glasses, and multiview lenticular display (MLD). Some of these are illustrated in Figure 1. Although HMDs are often used to provide an immersive and borderless 3D environment, as beneficial for gaming and simulation environments, these suffer a drawback of isolating individual users from the physical world. As this degrades group interaction, HMD technology is less suited to the educational or collaborative contexts.

Polarized stereo projection passes two images through orthogonal polarizing filters, superimposed onto the same screen. These are viewed through eyeglasses containing a matched pair of polarizing filters so that each eye sees a different image. In the two-color anaglyph technique, the two views are projected together as a single image, through different color filters, usually red and blue. Similarly, when viewed through a pair of matched color filters, each eye sees a different image, although with an obvious loss of color quality. LCD shutter glasses alternate blanking the view of each eye, synchronized to the display of an alternating left-right view. These technologies, however, can only render a single view for all participants, and often limit the viewers' position to a specific sweet spot. This can be problematic in the context of planning or training.

Multiview lenticular technologies employ an array of transparent lenses over a flat display surface that selects from different pixels (typically divided by column) based on viewing angle. This allows for experience of both stereopsis and motion parallax without restricting mobility or imposing the need to wear special eyegear. These displays offer a large viewing zone and are thus well-suited to simultaneous viewing by multiple individuals. Display techniques are discussed in further detail in Section 3.3.

2.2 Monoscopic vs. Stereoscopic Displays

Early studies investigated the comparative benefits of depth cues from a silk cursor and stereopsis[3] and the improved understanding of complex 3D data provided by stereoscopy.[4,5] In a precursor to the study described here, we found that participants were able to perform the same positioning task under stereographic viewing conditions in only 75% of the time and with half the error as in the monoscopic condition, given the positioning task within the volumetric medical data.[6]

Experiments with immersive stereoscopic visualization demonstrate that display of neurological data in 3D leads to an improved level of interpretation,[7] due in part to the availability of depth cues. More recent work[8] suggests the value of integrating multimodal data in the same visualization, as well as incorporating a representation of uncertainty in the display. However, these effects are not necessarily universal. A study of laparoscopic surgery visualization suggest that the lack of depth and spatial orientation cues do not significantly affect task performance.[9] Rather, individual experience was found to be the primary determinant of performance. Despite these results, the authors reported that two thirds of the test surgeons claimed that depth perception did help their performance.

More recent experiments found that while volumetric displays may be helpful for basic tasks involving judgement of depth or collisions, they do not, in their present form, aid in the comprehension of more complex 3D scenes.[10,11]

2.3 Medical Data Visualization

Traditionally, 3D medical data has been viewed by indirect volume visualization methods such as cutting planes or isosurfaces. In the former case, the results appear as 2D textures, and in the latter, as extracted intermediate geometric surfaces, rendered as polygonal meshes. For both, only a subset of the 3D data is used to generate the output image, for example, only selecting voxels on the cutting plane or those close in value to the isovalue. Both mesh discretization and sub-selection of data risk incurring loss of detail. In contrast, using direct volume rendering (DVR), the entire data set can potentially contribute to the output image.[12] This approach has proven valuable for many visualization tasks, including protein data visualization[13] and multi-million atom quantum dot simulation.[14] Although more computationally intensive, DVR is better suited to present structural details and the relationship between different materials, as illustrated in Figure 3. Furthermore, DVR settings can be adjusted to increase the opacity of regions of interest, or modify the transparency of surrounding volumes.

Other neurological visualization tools exist for viewing diffusion tensor MRI (DTI-MRI) with interactive volume rendering, suitable for volumetric vector-field representation,[15] in which users can interactively adjust the parameters to better understand relevant properties such as depth relationships. Tools have also been developed to permit the surgeon to switch between views of skin, cortex, tumor, vessels, and ventricles, including a set of automated scripts for cortex extraction, based on the Brain Extraction Tool (BET)*.[16] Some commercial systems, such as Medtronic StealthStation†, also integrate MRI slices with an overlaid view of the surgical instrument, obtained from real-time tracking data.

2.4 Interaction

Effective interaction with the visualization is also of great importance and depends a great deal on the choice of input device for manipulating the display. Direct manipulation tends to be more efficient than keyboard-mouse control,[17] assuming accurate, low-latency input. In this manner, a small number of simple hand gestures or

*http://www.fmrib.ox.ac.uk/fsl/

†http://www.medtronicnavigation.com/procedures/navigation/systems.jsp

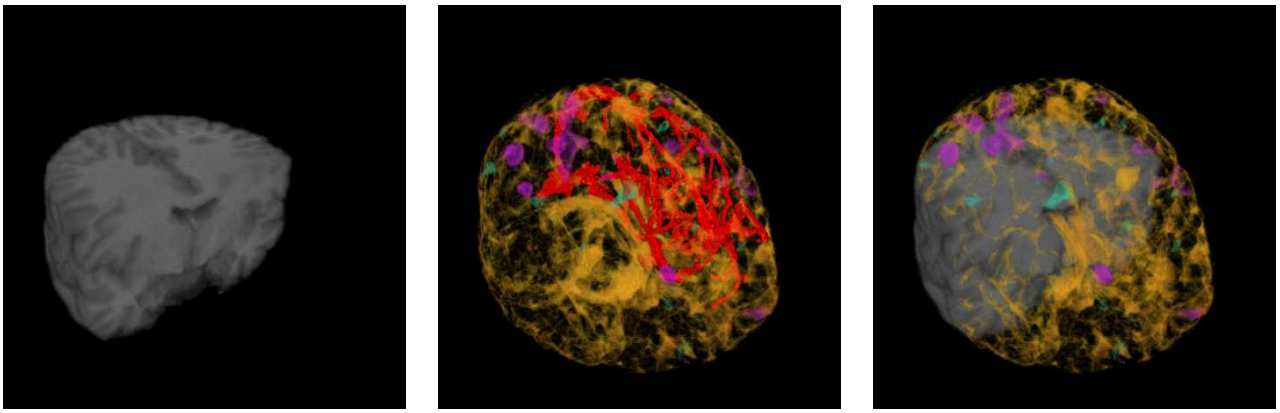

(a) (b) (c)

Figure 2. We illustrate the integrated visualization results: (a) cutting plane result; (b) direct volume rendering of mixed modalities, including blood vessels (red); (c) combined view. The combined view (c) shows the direct volume rendered MRI (yellow), with aligned fMRI (cyan), PET (magenta) data, and the cutting plane result (gray). Volumetric context and the combination of multi-modality data help surgeons to understand the medical data in a more effective way.

tangible objects allows for effective manipulation of a three-dimensional visualization, with less training than is required to achieve comparable control through a traditional graphical user interfaces (GUI).

Input devices one might consider range from keyboard, trackball, 6D controllers, stylus, or in the case of Handsaw,[18] to the position of an outstretched hand or a laser pointer.

Hinckley et al.[17] demonstrated the benefits of a Tangible User Interface (TUI), which perceptually couples a physical representation to its corresponding digital display. For example, the user can manipulate a doll's head, corresponding to the brain volume and a plastic plate for the cutting plane to determine the viewed cross-section. This approach offers a distinct advantage in supporting natural two-handed gestures for control. Although such a TUI may be ideal for planning or teaching tasks, it is ill-suited for use in an actual surgical procedures. In such a context, the surgeon's hands are often occupied, for example, working with surgical tools, and concerns of sterility are critical. We must thus ensure that similar functionality is available even without the affordances of physical representations for manipulation. For this purpose, we consider that supporting interaction with bare- or gloved-hand gestures may be imperative.

3. DESIGN APPROACH

3.1 Display of Multiple Data Sets

Although not yet experimentally validated, we expect that significant benefits may arise from the option to integrate any or all data in the same visualization, on demand. Working with neurosurgeons, we are attempting to characterize the parameters that maximize the value of information displayed during surgery. One of the most fundamental of these parameters is the choice of data sets (MRI, fMRI, PET, DTI, etc.) that best illustrate the relevant content, for example, blood vessels, functional regions, or white matter tracts. While a 3D view of the full volume is useful to understand spatial relationships in image-guided neurosurgery, a 2D cutting plane view might be preferred when detailed texture information needs to be analyzed. As shown in Figure 2, both a direct volume rendering and a traditional cutting plane view can be integrated in the same view simultaneously, thereby illustrating not only the detailed texture of the scanned data but also its relationship to data at other depths. The principal advantage of such integrated display is that all data sources can be viewed simultaneously, without the need to switch between them, and thus, cross-modal relationships can also be observed directly.

3.2 Integrated Direct Volume Rendering

We employ slice-based volume rendering to create high-quality DVR images, allowing for reasonable performance on commodity graphics hardware. The data are drawn on planes, which slice the voxels at a slanted angle perpendicular to the viewing direction, as shown in Figure 3(a) with a sample rendered result of an MRI scan of

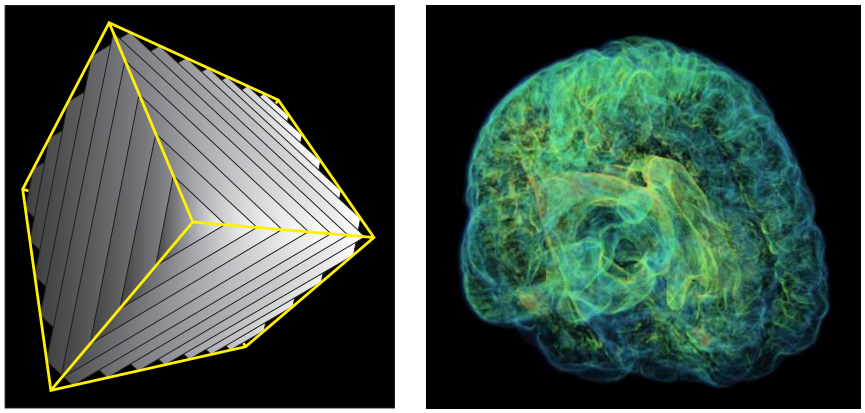

(a) slicing (b) rendering result

Figure 3. Slice-based direct volume rendering (DVR). 2D slices are sampled within the 3D volume (a), based on the viewing direction. Then these slices are drawn from back to front to achieve the direct volume rendering result (b).

the human brain in Figure 3(b). Importantly, these 2D slices are generated based on desired viewing direction, which may not align with the original axis along which the MRI images were obtained.

Each slice is a 2D texture, with texel components $(\mathbf{c_t}, \alpha_t)$, where $\mathbf{c_t}$ is the set of color components (r_t, g_t, b_t), and α_t is opacity. These components are interpolated within the 3D volume to form voxels with components $(\mathbf{c_v}, \alpha_v)$, where $\mathbf{c_v}$ is (r_v, g_v, b_v), calculated from from the contributions of all desired n data sources, each with density $D_i(x, y, z)$. As introduced in traditional ray casting algorithm,[12] each voxel is assigned a color $\mathbf{c_{vi}}$ and an opacity α_{vi}, according to the density D_i. The final color $\mathbf{c_v}$ and opacity α_v are the combinations of $\mathbf{c_{vi}}$ and α_{vi}, respectively:

$$\alpha_v = \begin{cases} 1 & \text{if voxel is on or below the cutting plane} \\ \sum_{i=1}^{n} \alpha_{vi} & \text{otherwise} \end{cases}$$

$$\mathbf{c_v} = \begin{cases} \text{the grayscale value, } D_{mri}(x, y, z) & \text{if voxel is on or below the cutting plane} \\ \dfrac{\sum_{i=1}^{n} \mathbf{c_{vi}} \cdot \alpha_{vi}}{\sum_{i=1}^{n} \alpha_{vi}} & \text{otherwise} \end{cases}$$

where D_{mri} is the voxel density of the MRI data on or below the chosen cutting plane. Note that the display of the cutting plane is only intended for the voxels within the brain region of the MRI data, which represents physiological structure of the soft tissue in the brain. When α_v is set to 1, the corresponding voxel is totally opaque. Example results are illustrated in Figure 2.

3.3 Stereoscopic Display

We have so far experimented with stereoscopic viewing using both a head-mounted display and a multiview lenticular display. For the head-mounted display, stereopsis is obtained by presenting different images on the left and right LCD screens seen respectively by each eye. This is achieved by enabling OpenGL quad-buffered stereo mode (`GLUT_STEREO`), and separately rendering the left and right views to their corresponding buffers, with the flags `GLUT_LEFT` and `GLUT_RIGHT`, respectively. Similar techniques can be used for polarized stereo or LCD shutter glasses.

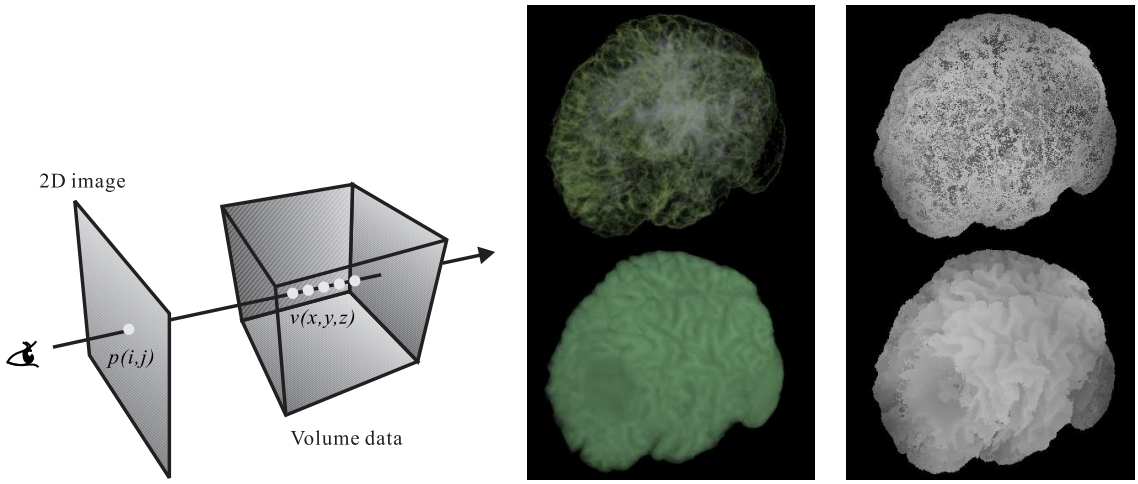

(a) direct volume rendering (b) 2D image (c) depth map

Figure 4. Depth approximation for direct volume rendering. Since the volume is usually translucent, the final color of pixel $p(i, j)$ is based on contributions from a set of voxels $v(x, y, z)$, as shown in (a). The depth value of this pixel is thus not uniquely defined. Two examples are provided where the same MRI data is rendered with different transfer functions to generate a translucent volume (upper row) and an opaque volume (lower row). Here, we show (b) the 2D images, and (c) the approximated depth maps. The resulting depth map appears perceptually acceptable for largely opaque data but is clearly a poor approximation otherwise. A depth-map value of zero corresponds to a distance of infinity.

While some multiview lenticular displays (MLD) take as input multiple pre-rendered views, certain models[‡] require a regular 2D texture representation (Figure 4(b)) and its corresponding depth map (Figure 4(c)), which indicates the distance from the eye to each visible point of the image. From these two representations, the graphics hardware synthesizes multiple views, corresponding to different viewing angles. When observed from an appropriate distance, two different views are presented to the left and right eyes, resulting in the appearance of depth in the scene.

Unfortunately, for direct volume rendering of semi-translucent data, there are often many pixels that are not characterized by a single depth value, as contributions to each pixel may be combined from a set of *related* voxels at varying depths. Assume the set of related voxels $v(x, y, z)$ contribute to pixel $p(i, j)$. As seen in Figure 4, if we select the voxel with maximum opacity to represent the depth of that pixel, the resulting depth map appears perceptually acceptable for largely opaque data but is clearly a poor approximation otherwise. Until a better solution can be found, we expect that stereoscopic rendering based on individually rendered left/right views will be more effective.

3.4 User Interaction

To explore the utility of different interaction paradigms, we implemented both a tangible user interface (TUI), inspired by the work of Hinckley et al.,[17] and a free-handed gestural control, as warranted for the actual surgical context.

In the absence of physical constraints, our preference is to manipulate the volume with one hand and perform additional functions (e.g., adjust the cutting plane) with a second. We presently support rotation and zoom of the volume, orthogonal and arbitrary cutting plane control, selective integration of different data sources, and simple deformations of the volume. More complex interactions could also be supported at the cost of increased training time. These may be included in our software as deemed beneficial by our intended users.

3.4.1 Tangible User Interface

Our initial TUI consists of three physical objects, shown in Figure 5: a ball representing the volume, a cube for control of the orthogonal cutting plane, and a rigid plane (a CD jewel case) to control the position and orientation of the arbitrary cutting plane. For prototyping purposes, these objects are tracked using IR-reflective markers.

[‡]For example, the Philips 3D WOWvx series.

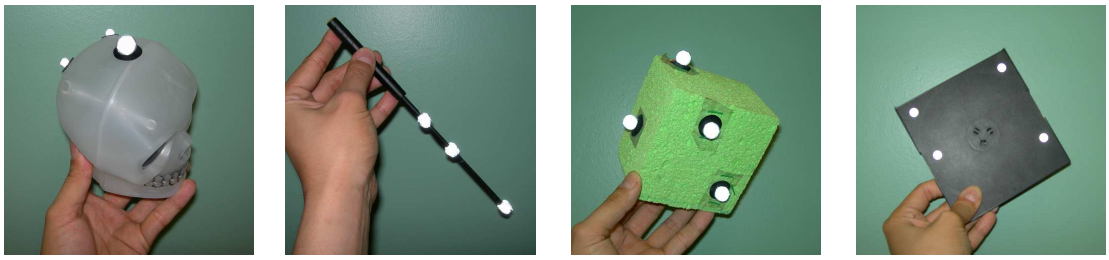

(a) grab (b) probe (c) orthogonal cutting (d) arbitrary cutting

Figure 5. Tangible objects: (a) a ball to grab the volume, (b) a replica of a surgical probe to manipulate the virtual probe, (c) a cube to control the orthogonal cutting plane, and (d) a rigid board to change the arbitrary cutting plane.

(a) glove (b) makers (c) grab (d) orthogonal cutting (e) arbitrary cutting

Figure 6. Hand gestures. (a) glove with IR-reflective markers for tracking. (b) organization of markers and segments. There are 12 markers ($1 \sim 12$) and 6 segments ($I \sim VI$) in total: two markers per finger, (root and joint) and two markers (11 and 12) for tracking the palm. This organization helps maintain a fixed length for each segment. We currently implement three different gestures: (c) grab, (d) orthogonal cutting plane manipulation, and (e) arbitrary cutting plane control.

3.4.2 Hand Gesture Interface

For free-handed interaction, we employ a simple gestural vocabulary, shown in Figure 6, consisting of "grab" (pinching all fingers) for rotation and zoom operations, pinching thumb and index finger for orthogonal cutting plane manipulation, and an open hand for arbitrary cutting plane control. As noted earlier, we expect at least one of the surgeon's hands to be occupied during large portions of the operation, holding surgical instruments. Thus, despite the obvious benefits of bimanual interaction, we attempt to realize the principal functionality of the visualization system using only single-handed gestures.

Hand gestures are identified by the relative positions of a set of twelve markers ($1 \sim 12$), two of which define the palm and the remaining ten correspond to the root and joint of each finger, as shown in Figure 6(a). Each of the six segments ($I \sim VI$), shown in Figure 6(b), connects a pair of markers, with an assumption that the distance between these remains constant.

4. EXPERIMENTAL DESIGN

Returning to our objective of evaluating the efficacy of different displays in the context of a medical visualization task, we compared the stereoscopic HMD (800x600 pixel resolution), with polarized projection (1024x768 resolution), and multiview lenticular display (MLD) (1920x1080 absolute resolution, but 960x540 effective resolution).

Twelve participants took part in the experiment, 6 male and 6 female, ranging in age between 23 and 36. The experimental task was to define a straight vessel-free path from the cortical surface to the targeted tumor using (a pen-like) probe. For this task, we utilized a geometric model of the segmented brain vasculature, as shown in Figure 7(a), instead of the volumetric representation. This simplification allowed us to render the stereoscopic view on our MLD, which requires texture plus depth as input, as discussed in Section 3.3.[§] It is also worth noting that MLDs typically segment the allowable viewing angle into smaller zones, each of which repeats the same content. Users must remain within a single zone in order to experience a correct stereoscopic effect.

[§]As noted above, this prevents us from achieving an effective rendering of semi-transparent data on the MLD.

(a) (b)

Figure 7. (a) sample view of the experimental configuration, including the vasculature (red), virtual probe (silver) and tumor (green). (b) experimental setup including a plastic skull, a replica of the surgical probe, a head-mounted display, and a dance pad.

Subjects were provided with hand-coupled motion cues as they manipulated the volume and probe (Figure 7(b)). Two tangible objects were used as input devices: a small plastic skull to manipulate the orientation and proximity (zoom) of the vasculature, and a replica of a biopsy probe used to locate the tumor. Once the subject was satisfied with the probe position, a foot step, sensed by an electronic game controller dance pad, is used to end the trial. This avoided the risk of introducing jitter to the probe position, which may have resulted otherwise had a manual button press been required.

The order of display conditions was scheduled in a Latin square design across both male and female subjects. Participants ran 30 trials in each condition, consisting of 10 easy, 10 medium, and 10 difficult-to-reach tumor positions, presented in random order. The same trials were then repeated in the subsequent display condition, again, with the tumor positions presented in random order.

5. RESULTS AND DISCUSSION

For each trial, we measured the average completion time and the average position error, computed as the distance between the probe tip and the center of the tumor. We also measured the amount of rotation of the (TUI) skull and the translation of the probe, to compare how viewing condition affects the amount of manipulation of the display. The results shown in Figure 8(a) and (b) indicate that overall, polarized projection outperformed the other two display modalities, exhibiting both the fastest completion time and the smallest error. p values show that they are all significantly different. HMD performance was slightly inferior and the MLD significantly inferior. This suggests a strong correlation between resolution and task performance, which will need to be verified by repeating the experiment with the polarized projection at the same resolution as the HMD. Details of statistical p values are given in Table 1.

	overall	MLD vs. HMD	MLD vs. Polarize	HMD vs. Polarize
Time	$p < 0.05$	$p < 0.01$	$p < 0.01$	$p < 0.05$
Error	$p < 0.01$	$p < 0.01$	$p < 0.01$	$p < 0.01$
Skull's rotation	$p > 0.05$	$p > 0.05$	$p > 0.05$	$p > 0.05$
Probe's translation	$p < 0.05$	$p < 0.01$	$p < 0.05$	$p < 0.01$

Table 1. Statistically significantly different variables across all participants

Manipulation of both the skull and probe (significantly different) was also affected by the display (Figure 8(c) and (d)). Interestingly, use of the HMD led to the greatest amount of manipulations, possibly because this display prevents the participant from seeing the TUI while manipulating the input devices. With regard to gender differences, we observed that, with HMD and polarized projection, male subjects were significantly faster than female ($p < 0.05$). On average, female subjects performed more accurately with the MLD ($p < 0.01$), but less accurately with the HMD ($p < 0.05$), than male subjects.

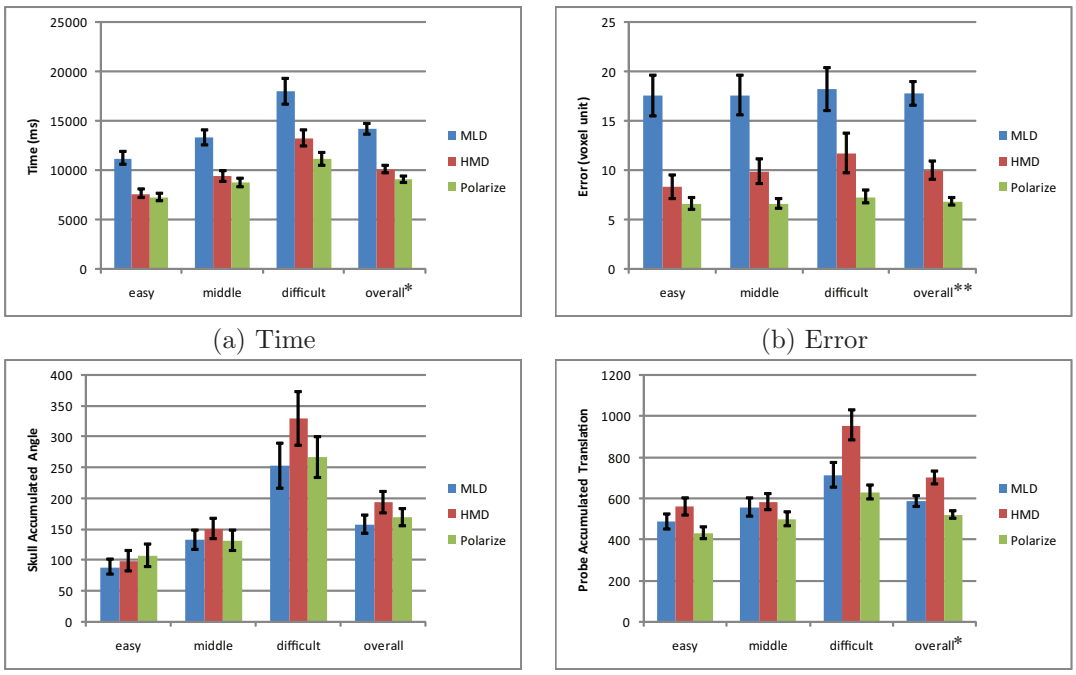

(a) Time (b) Error

(c) Skull's rotation (d) Probe's translation

Figure 8. Experimental results. The first three groups of bars plotted relate to the level of difficulty to localize the tumor and the fourth pair are averages across all difficulty levels. Standard error is indicated by the error bars. '$*$' denotes $p < 0.05$, and '$**$' denotes $p < 0.01$.

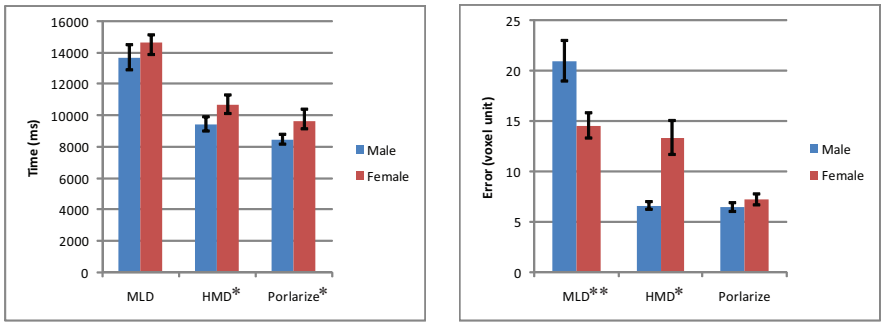

(a) Time (b) Error

Figure 9. Gender differences. '$*$' denotes $p < 0.05$, and '$**$' denotes $p < 0.01$.

In a previous experiment, comparing monoscopic and stereoscopic viewing conditions, we also measured the minimum distance between probe and vasculature at the conclusion of each trial. Since no significant difference was observed, this measurement was not included in the later experiments.

The quantitative measurements were reinforced by the subjects' responses to our questionnaire, as summarized in Figure 10. Participants were asked to rate the three displays on a 5-1 Likert scale (5 being the strongest) according to (a) "How well did you feel that each display conveyed depth information?"; (b) "How comfortable was it to work with each display?"; (c) "How good was the image quality with each display?"; (d) "Overall, how would you rate each display?". Both polarized projection and HMD were judged to convey depth information effectively, the HMD was the least comfortable, and the MLD presented the worst image quality. Overall polarized projection was given the highest rate, consistent with the task performance measures described earlier.

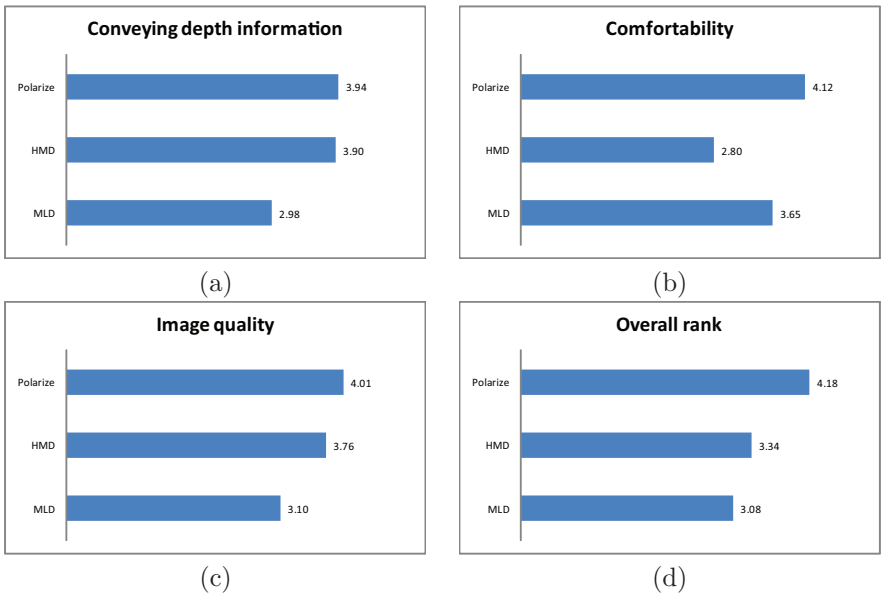

Figure 10. Average ratings of each display on a 5-1 Likert scale, 5 being the strongest.

6. CONCLUSIONS AND FUTURE WORK

We have described an immersive system for medical data visualization, with applications to neurosurgical training, planning, and practice. In addition, we discussed an experiment involving a comparison of three stereoscopic display modes. We observed performance advantages of polarized projection display, likely due to both high image quality and comfort. Future experiments may include comparative analysis with other 3D technologies such as swept-surface volumetric displays.

Our objectives, throughout the development of our system and the experiments described here, are the improvement of task performance and augmentation of the surgical training environment for more effective communication of experience and expertise. Although our experiments were restricted to the use of hand-held interface devices, which limit the applicability of the system to training and planning scenarios rather than actual surgery, we are also investigating the use of non-contact 3D gestures to achieve similar tasks. We believe that there is a tremendous potential to improve efficiency of neurosurgical procedures by offering the surgeon the capability of direct manipulation of a data visualization during surgery.

Other extensions include the use of composite gestures,[19] employing both static and dynamic components, and the possible addition of speech recognition to invoke modes that are less easily mapped to motor actions. Speech input may also be valuable in situations where a simple verbal utterance can describe a desired change to the display, thus supporting interaction requirements when the surgeon's hands are occupied.

Given the potentially large size of data sets, it may prove necessary to consider an adaptive processing strategy to ensure adequate update rates, similar to the manner in which Zhang et al.[7] employed a decimation algorithm to reduce the complexity of the models for DT-MRI volume visualization. Current MRI data tend to be limited in resolution (e.g., the Simulated Brain Database[¶] is sampled at $181 \times 217 \times 181$ voxels). However, availability of higher resolution content, such as atlas data or microscopy scans in the order of one billion pixels or greater, motivate consideration of a hierarchical volumetric representation for more efficient rendering.

7. ACKNOWLEDGEMENTS

The authors wish to acknowledge the generous support of Fonds québécois de la recherche sur la nature et les technologies (FQRNT) and the Natural Science and Engineering Research Council (NSERC), who have funded the research described here. The authors thank Adriana Olmos, François Bérard, Alvin Law, and Jessica Ip for

[¶]http://www.bic.mni.mcgill.ca/brainweb/

the experiment setup. In addition, the authors are most grateful for the assistance in the statistical analysis provided by Mitchel Benovoy and the valuable feedback offered by colleagues Tal Arbel, Louis Collins, Laurence Mercier, Rolando del Maestro, Kevin Smith, and Mike Wozniewski, who have all provided important suggestions that influenced the evolution of this work.

REFERENCES

[1] Wössner, U., Schulze, J., Walz, S., and Lang, U., "Evaluation of a collaborative volume rendering application in a distributed virtual environment," in [*EGVE '02: Proceedings of the workshop on Virtual environments 2002*], 113–122 (2002).

[2] Guan, C., Serra, L., Kockro, R., Hern, N., Nowinski, W., and Chan, C., "Volume-based tumor neurosurgery planning in the virtual workbench," in [*Virtual Reality Annual International Symposium*], 167–173, IEEE (March 1998).

[3] Zhai, S., Buxton, W., and Milgram, P., "The "silk cursor": investigating transparency for 3d target acquisition," in [*CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*], 459–464 (1994).

[4] Arthur, K. W., Booth, K. S., and Ware, C., "Evaluating 3d task performance for fish tank virtual worlds," *ACM Transactions on Information Systems* **11**(3), 239–265 (1993).

[5] Ware, C. and Franck, G., "Evaluating stereo and motion cues for visualizing information nets in three dimensions," *ACM Transactions on Graphics* **15**(2), 121–140 (1996).

[6] Wang, G., Mercier, L., Collins, D. L., and Cooperstock, J. R., "A comparative study of stereoscopic displays for identication of a vessel-free path," in [*Medicine Meets Virtual Reality 2009 (MMVR17)*], (February 2009).

[7] Zhang, S., Demiralp, C., Keefe, D., da Silva, M., Laidlaw, D., Greenberg, B., Basser, P., Chiocca, E., Pierpaoli, C., and Deisboeck, T., "An immersive virtual environment for DT-MRI volume visualization applications: A case study," in [*Proceedings IEEE Visualization 2001*], 437–584 (October 2001).

[8] Simpson, A. L., Ma, B., Chen, E. C. S., Ellis, R. E., and Stewart, A. J., "Using registration uncertainty visualization in a user study of a simple surgical task," in [*MICCAI (2)*], 397–404 (2006).

[9] Chan, A. C. W., Chung, S. C. S., Yim, A. P. C., Lau, J. Y. W., Ng, E. K. W., and Li, A. K. C., "Comparison of two-dimensional vs three-dimensional camera systems in laparoscopic surgery," *Journal Surgical Endoscopy* **11**, 438–440 (May 1997).

[10] Grossman, T. and Balakrishnan, R., "An evaluation of depth perception on volumetric displays," in [*AVI '06: Proceedings of the working conference on Advanced visual interfaces*], 193–200 (2006).

[11] Grossman, T. and Balakrishnan, R., "The design and evaluation of selection techniques for 3d volumetric displays," in [*UIST '06: Proceedings of the 19th annual ACM symposium on User interface software and technology*], 3–12 (2006).

[12] Levoy, M., "Display of surfaces from volume data," *IEEE Transactions on Visualization and Computer Graphics* **8**, 29–37 (May 1988).

[13] Hu, M., Chen, W., Zhang, T., and Peng, Q., "Direct volume rendering of volumetric protein data," in [*Computer Graphics International 2006*], 397–403 (2006).

[14] Qiao, W., Ebert, D. S., Entezari, A., Korkusinski, M., and Klimeck, G., "VolQD: Direct volume rendering of multi-million atom quantum dot simulations," in [*IEEE Visualization 2005*], 319–326 (2005).

[15] Wenger, A., Keefe, D., Zhang, S., and Laidlaw, D., "Interactive volume rendering of thin thread structures within multivalued scientific datasets," *IEEE Transactions on Visualization and Computer Graphics* **10**, 664–672 (November/December 2004).

[16] Smith, S., "Fast robust automated brain extraction," *Human Brain Mapping* **17**, 143–155 (November 2002).

[17] Hinckley, K., Pausch, R., Proffitt, D., and Kassell, N., "Two-handed virtual manipulation," *ACM Trans. Comput.-Hum. Interact.* **5**(3), 260–302 (1998).

[18] Bonanni, L., Alonso, J., Chao, N., Vargas, G., and Ishii, H., "Handsaw: tangible exploration of volumetric data by direct cut-plane projection," in [*CHI 2008*], 251–254 (2008).

[19] Corso, J., Ye, G., and Hager, G., "Analysis of Composite Gestures with a Coherent Probabilistic Graphical Model," *Virtual Reality* **8**(4), 242–252 (2005).