Multiple Views and Magic Mirrors - fMRI Visualization of the Human Brain

Category: research

Abstract

Multimodal visualization of functional and anatomical data of the human brain is an important field in medical volume visualization. The aim of this application is to provide the user with information on the location of functional activations in the different regions of the brain. When the approach of direct volume rendering is chosen, the visual impression of details usually suffers from accumulating the image from colors and opacities derived from the data set. We present extensions to overcome this problem: *Transfer function volumes* are used for the highlighting of activated regions. *Multiple views* simplify the task of localizing these parts of the brain. Complementary information is brought into the visualization by *Magic Mirrors* in order to enhance the comprehensive view of the multimodal arrangement of volume data sets.

Keywords: Volume Visualization, Multimodality, fMRI

1 Introduction

Over the last decade, the use of digital modalities for clinical diagnosis and therapy planning has dramatically increased. New developments in three-dimensional volume acquisition techniques such as Spiral CT (*Computed Tomography*) scanning, faster and better pulse sequences for MRI (*Magnetic Resonance Imaging*) and improved functional imaging provided by modalities such as PET (*Positron Emission Tomography*), SPECT (*Single Photon Emission Computed Tomograph*), EEG (*Electro EncephaloGraphy*), MEG (*Magnetic EncephaloGraphy*), and fMRI (*functional Magnetic Resonance Imaging*) are creating a rapidly increasing demand for integrated multimodality 3D visualization. It aims at efficiently presenting information from multiple modalities, in order to obtain a comprehensive picture of clinical cases. Therefore, depending on the clinical requirements, it might often not be sufficient to consider vascular, anatomical and functional information separately but to superimpose images from different modalities.

Functional MR (magnetic resonance) imaging is used to determine regions of the human brain, which are responding to sensory stimulation, motor or cognitive activity or emotional states. Neuroimaging experts investigate such scans usually by employing their talent to visually extract the characteristics of activated regions from the scanned volumes. This exploration task is usually done on two-dimensional slices of the human brain, where the functionally activated regions are highlighted in some special way (see Fig. 1, top row). Viewing several slices at once (side by side) is referred to as the *image galleries* approach. This way of investigating the data sets provides detailed 2D anatomical context information. Nevertheless, this way of dealing with multiple data sets does not aid the user in forming a three-dimensional model of the underlying spatial relationship of functional activations and the anatomical geometry. Other approaches try to increase the amount of presented information by utilizing *maximum intensity projections* (MIP), mapped onto 2D reference grids with brain outlines [3] (see Fig. 1, middle and bottom row). This representations aid the task of side by side comparisons, but nevertheless provide just very rough information of the anatomical context. Therefore it is very hard for the user to mentally reconstruct the three-dimensional structure of the patients brain.

Three-dimensional surface renderings of the brain, where the functional activation is encoded by coloring the cerebral cortex provides a misleading anatomical context. Activated subcortical foci below the surface of the brain are obscured. Furthermore, all foci seem to lie on the surface of the cortex, as no depth cues are provided.

Direct volume rendering would be able to overcome this problem. Assigning transparencies to data ranges would allow the deeper structures to be visualized. Therefore direct volume rendering would be the most usable approach to the solution of this problem. Usually, however, anatomical detail has to be sacrificed, when employing transparencies [4], decreasing the visual appearance of details significantly.

For these reasons K. Rehm and other authors recently proposed, that the anatomy of the brain should not be visualized three-dimensionally at all [14]. They use symbolic representations known from other fields of visualization (like fluid or information visualization) to encode the location and size of foci. In order to place the foci into anatomical context, they are projected onto the walls of a cube like environment together with different representations of the brain. We are going to utilize a similar approach of adding information to the visualization in Section 4.

Direct volume rendering still seems to be the most useful approach to clearly give visible hints on the location and size of activated regions. We are going to present some extensions to overcome the problem of reduced visibility in the following sections.

First of all, the given problem of our special case of

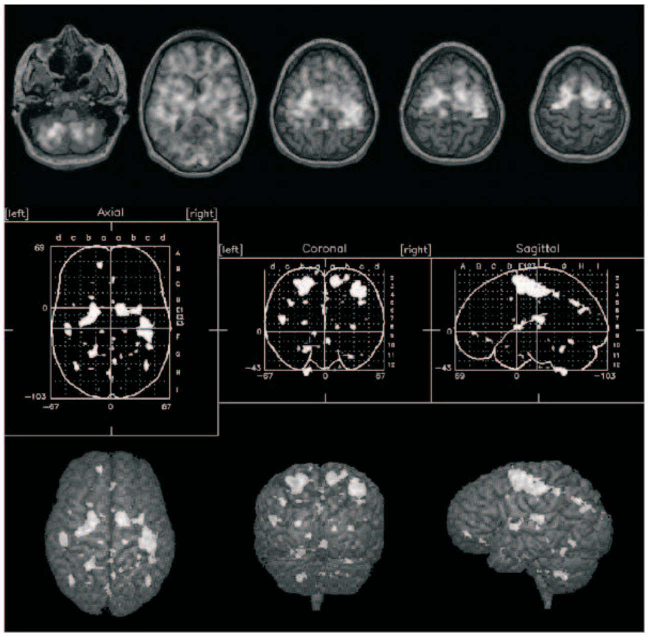

Figure 1: Conventional approaches to multimodal visualization [14]. Top, selected MRI slices are merged with functional data by pixel interleaving [15]. Middle, the maximum intensities of the data are projected onto axial, coronal and sagital planes bearing brain outlines and millimetric Talairach coordinates [18]. Bottom, the maximum data intensities are projected along the line of sight onto the cortical surface.

combined visualization of functional and anatomical data of the human brain will be stated in detail. The problems arising when using standard direct volume visualization techniques will be discussed in Section 3. The first improvement using multiple views will be discussed in Section 4. Additional aids for the investigating user (Magic Mirrors) are introduced in Section 5. A results section will provide images and statistics to prove the usability of our techniques. Finally conclusions will be drawn.

2 The problems encountered in multimodal volume visualization

In this work a volume visualization technique for multimodal data of the human brain is presented, consisting of functional data of the human brain (measured by fMRI) and anatomical data of the brain (measured by different methods of MRI). Other types of multimodality are out of the scope of this paper and are not discussed. The functional data represents the activation potential of the brain at a certain point in time.

The test situation used was the following: A patient has to perform the repetitive action of thumb and index finger tapping. The activation potential of the brain is measured. This gives the functional context of the reaction to the action of the fingers. The fMRI scans are acquired during stimulation as well as under resting conditions. By calculating the differences of these images the activated regions of the brain can be obtained. A review of the physics of MRI (and fMRI) resides outside of the scope of this paper. For a complete description of MRI specialized literature [2, 17] should be consulted.

The problems encountered with multimodal volume visualization of anatomical/functional data of the human brain include the following:

- *the data sets have different resolutions*: the resolution of the functional data sets is usually significantly lower than that provided by the anatomical data. This is due to the fact, that the functional data has to be scanned interactively. An anatomical MRI scan is acquired using more time.

- *the spatial extend of the data sets is different*: again, due to the time constrains in fMRI measuring. Usually only parts of the brain (for example only 5-7 slices) are scanned for the functional data set.

Therefore the data sets first have to be spatially aligned through *registration* [12]. The actual registration process for our test data sets happened to be rather trivial. The spatial locations of the two data sets in relation to each other were known exactly due to the information provided by the scanning modality. MR-typical distortions due to magnetic field inhomogenities were negligible. Only the different resolutions and spatial extends of the two data sets had to be taken into account.

Afterwards visualization is used to identify the regions of the brain, where the changes in the activation potential

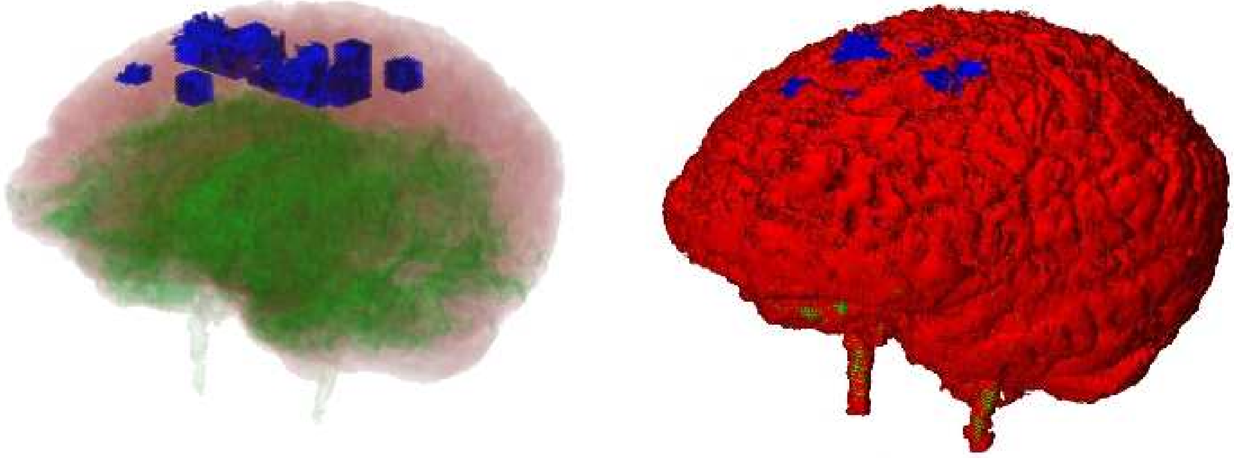

Figure 2: Different problems arise when using direct volume rendering. When surface rendering is employed, the regions of activation are hidden by the cerebral cortex (right hand side). Left, very low opacity allows deeper structures to be visualized. However, it is hard to tell, where the foci reside exactly.

occurred. This visualization step is the main point of the presented paper.

3 Transfer-function volumes for multimodal visualization

As mentioned above, the usage of shaded surface displays (like for instance *isosurface extraction by marching cubes* [11]) does not serve our purpose of visualizing multimodal data very well. First of all, inner structures of the brain (like white matter) are obscured by the covering surface. When activated regions are encoded by coloring the according parts of the surface, incorrect information is presented: activations seem to appear only on the cortical surface. The parts within the brain responding to activation are not visible at all. The usage of transparent surfaces could overcome this problem. Problems with the reduced visual appearance of highly transparent surfaces could also be compensated by special strategies [16, 5, 8, 6, 7]. Nevertheless, due to the very low resolution and quality of the functional MRI data sets, it is not possible to define the exact border region of the activated foci as it would be necessary for the extraction of isosurfaces. Therefore the authors utilized Direct Volume Rendering (or probabilistic rendering) techniques for their approach, which are concerned with the visualization of fuzzy, cloud-like structures. As data values are mapped into color and opacity values by the usage of transfer functions, during image rendering the physical properties of light transportation in non-opaque media can be simulated. This approach serves the given problem of multimodal volume visualization very well.

```
∀ voxels (x,y,z) ∈ MRI data set
  Set TFV(x,y,z) to normal
Interpret registration information
∀ voxels (a,b,c) ≠ 0 ∈ fMRI data
(r,s,t)=Register.fMRI2MRI(a,b,c)
  Set TFV(r,s,t) to highlighted
```

Figure 3: Setting up the Transfer Function Volume

Among the variety of existing direct volume rendering techniques, *raycasting* [10] has been chosen for our approach to the visualization of fMRI data. As this method approximates techniques used in photorealistic rendering (global illumination), the simulation of physically correct phenomena like refraction or reflection is possible. Having a special enhancement to our technique in mind (described in Section 4), raycasting is the technique of our choice. A special scheme for the specification of transfer functions had to be employed. *Transfer-function volumes* are used to specify different transfer functions for separate regions of the data set to be visualized. A transfer-function volume is a three-dimensional array storing an index to a set of transfer functions for every single voxel. This index is used to access a transfer function for mapping the data value of the voxel into the physical properties color and opacity. When the number of transfer functions to be used is limited, e.g. to 4 or 8, the indices can be packed more densely and significant amounts of storage space are saved.

Assigning transfer functions on a voxel per voxel basis has several advantages:

- A special transfer function is defined, which should be used to encode the activated regions of the anatomy. During the *registration* process, the index of this transfer function is assigned to all voxels which are at least partially intersected by a voxel of the functional data set with activation potential. During rendering of the anatomical data set activated regions are colored using this special transfer function and are highlighted. Figure 3 gives a description of this process using a pseudo-code representation.

- Using a "NULL" transfer function (mapping every data value to zero opacity) *volume sculpting* facilities can be realized quite easily. Setting regions of the transfer function volume to the index of this special transfer function, various "cut-away" possibilities can be realized, because voxels values mapped by "NULL" transfer functions simply do not appear in the visualization. These regions can be specified by cutting planes, block cut operators, or even on a voxel by voxel basis. The anatomical data set used in the scope of this work has been segmented in this way. A segmentation mask gained by manual segmentation of the anatomical data set specified the voxels of the non-brain regions of the head scan, which were rendered using a "NULL" transfer function. The original data set has not been modified in any way.

- When this approach has to be optimized for the calculation time of the raycasting step, the data volume can be converted into a color/opacity volume using the transfer-function volume in a preprocessing step. If rapid changes of the used transfer functions are desired (for instance when rendering several frames for an animated visualization), the transfer-function volume does not have to be changed at all. Simply the set of transfer functions referenced by the indices is modified.

Direct volume rendering is capable to overcome the problem of occlusion. Nevertheless, significant detail in image quality is sacrificed when employing low opacity contributions of the voxels. Figure 2 shows a comparison of the two extremal cases. When a transfer function mapping the voxel values of the anatomical data set into high opacity values is used, the visualization yields to a surface like representation (right hand side). When very low opacity values are used, hardly any anatomical context is visible (left hand side). It is not easy to describe the location of activated regions according to the visualizations in Fig. 2. A solution to overcome this problem is shown in the next section. To compensate for the loss of image details, additional information is brought into the visualization by enabling alternative views of the brain.

4 Multiple Views - Enhancing the visual impression of direct volume renderings

From a single still image it is very hard to get the correct visual impression of depth. In the presented technique the image produced with the transfer-function volume is enhanced by adding three different views of the object to the final image. These multiple views are realized by positioning three image planes behind and below the object. A similar approach for a different task has been chosen by Mulder and van Wijk in 1997 [13]. They proved, that the visibility of features is greatly enhanced by these additional views.

Images displayed on this image planes are again generated by using the ray-casting technique [10]. Orthogonal projection has been proven to be perceptually superior to perspective projection in this application. Perspective distortion does not aid the user in interpreting the localization of features. This decision has also been made after the investigation of several CAD and modeling software packages using non-perspective top, left, and front views of the scene in addition to an perspective view. In these systems the views are simply arranged side by side, tiling the main window in an two by two arrangement. It is mandatory for the user to permanently have in mind, which viewing direction belongs to the single views. This situation is improved by arranging the views in a different way. A bit of detail and screen area of the views is traded for an easier understanding of their location. The three image planes are added to the scene as shown in Fig. 4. The white parts on the three planes (see Fig. 4, left) depict the location, where the additional images are displayed as textures (Fig. 4, right). Each of these images has it's own set of camera parameters. In the example of Fig. 4 the viewpoints for the image planes are behind (below, respectively) the planes. These image planes provide multiple views and are quite different in their properties from, e.g., mirrors or shadow walls. Placing a real mirror at the same location (an approach used by A. Kaufman [1]) would result in completely different images. Nevertheless, this different behaviour was aimed at for the further developments discussed in the next section.

The benefit gained from these multiple views is manifold:

- The location of activated regions becomes obvious even in still images, as the three-dimensional context of the two data sets can be recognized more easily when multiple views on the same scene are available. This is especially the case for highly transparent renderings (compare the image on the left hand side of Fig. 2 with Fig. 5, left hand side).

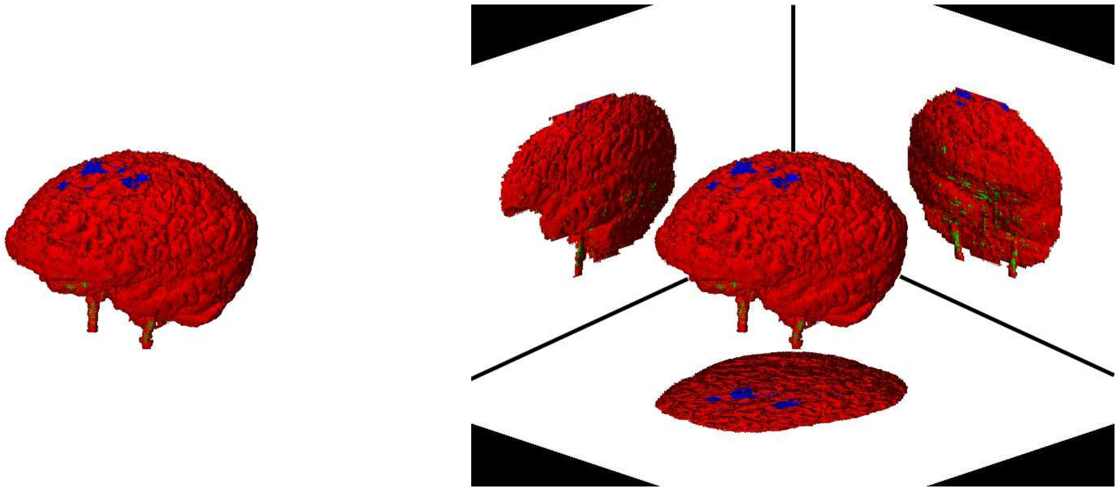

Figure 4: The principle of placing image planes as multiple views of the scene for enhancing volume visualization: without (left) and with (right) textures added.

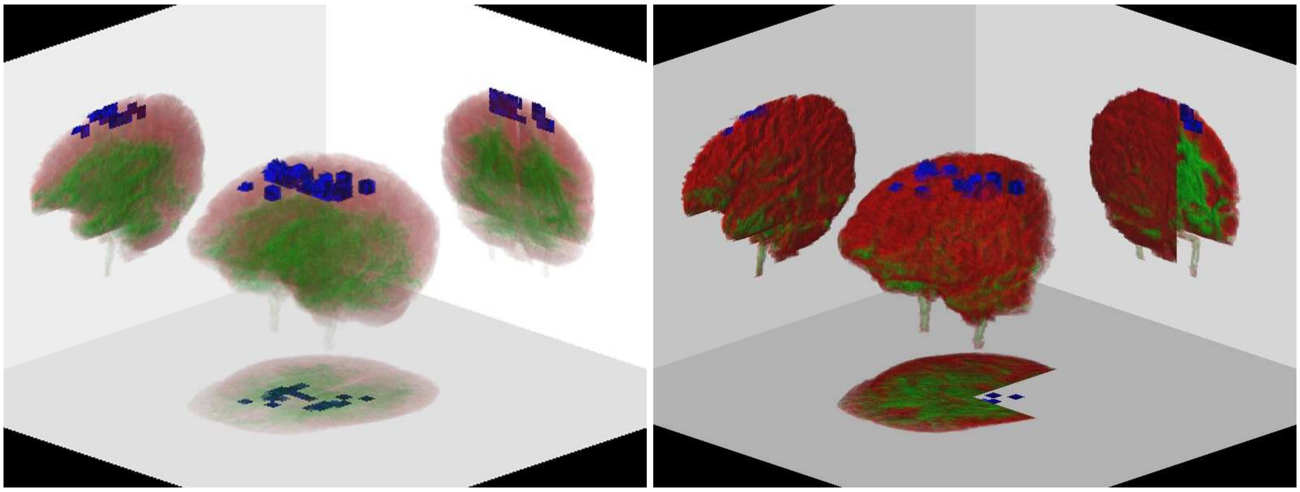

Figure 5: Left: Using multiple views, the exact localization of activations is also possible with transparent transfer functions; Right: Even occluded regions like the cutaway part become obvious with the help of multiple views

```
For dir∈ {top,left,front}
  tex[dir]=Render(MRI data, TFV, dir)
RenderMV(MRI data, TFV, tex)
```

Figure 6: Rendering a Multiple Views image: In a first step three textures are rendered, which are then mapped onto the mirrors during the final rendering step. The Transfer Function Volume created during registration is used for the visualization.

- The problem of occlusion is solved rather satisfactorily. In Fig. 5, right hand side, for example a block of the anatomical data has been cut out of the backside, but due to the multiple views provided by this technique, the locations of activations are still visible.

5 Magic Mirrors - Illustrating additional information

The benefits of multiple views of the brain generated by volume rendering can even be enhanced. As the employed image planes of the multiple views do not show the real world physical behaviour of , e.g., reflecting planes, several quite different rendering techniques can be used to generate the textures. Leaving the domain of physical reflection completely, the extended version of the multiple views approach has been named *Magic Mirrors*.

- *"Projection"* vs. *"Reflection"*: The multiple views textures used in the last section were created by raycasting the brain anatomy with an view point located behind (below, respectively) the image planes. For this reason, the quasi-reflective property of the views was gained. It is easy to locate the view point for texture generation on the other side of the brain geometry. Utilizing this approach, projections of the brain in direction of the mirror can be generated. The main views onto the brain (front, top, right) can be used to enhance the visual impression of the three-dimensional representation of the brain (Fig. 7, bottom mirror).

- Different transfer functions can be used for each of the Magic Mirrors. Various characteristics of the multimodal scenario can be emphasized in a single image. Fig. 7 demonstrates the relation of activated regions to white and gray brain tissue.

- Additional views to the anatomy created by non-direct volume rendering techniques can be used to provide complementary information. In Fig. 8 a 3D *"contour map"* of the brain and activated regions (generated by the ray-tracing approach) was mapped onto the Magic Mirror at the left hand side of the scene. The usage of other techniques like *maximum intensity projections*, *orthonormal cuts*, *color encoded cuts*, or *segmentation maps* as Magic Mirrors is also possible.

	anatomical	**functional**
slice resolution	256 x 256 pixels	64 x 64 pixels
slice count	64 slices	5 slices
voxel size	1 x 1 x 2mm	4 x 4 x 4mm

Table 1: Resolutions of the data sets used to get the results shown in this paper

6 Results

The goal of the presented work was to develop a technique, which allows the identification of the location of regions with activation potential. In order to prove the usefulness of the presented Magic Mirrors approach a non-clinical test scenario was set up. Still images as well as animation sequences of visualizations produced with the presented approach were to be evaluated.

The data sets used for this work were gained from a 3-Tesla MRI scanner. The anatomical data sets were acquired using a pulse sequence called *MDEFT* (Modified Driven Equilibrium Fourier Transform). These images are T1-weighted. The functional data sets used in this work are $T2^*$-weighted, and have already been pre-processed, so that the regions of activations are emphasized in the volume, and all other regions are neglected. The data sets featured the resolutions shown in Table 1. A small group of test persons was asked to interpret the given Magic Mirrors images (and animations) and verbosely describe the location as well as the size of activated regions in relation to the brain anatomy. All evaluators reported, that it was an easy task to derive the location of activated regions from the multiple views and Magic Mirrors representations. Especially the pre-rendered animation sequences simplified the users task to locate the foci.

The volume of the activated regions was not estimated as easily. Although this information is implicitly know to the system from the properties of the functional data set, it might be desirable to get an instant impression of the size of the activated volume. We are going to deal with this problem in our future work (see Section 7).

Electronic versions of the presented images as well as animations can be found at our web-site:
`http://www.cg.tuwien.ac.at/research/vis/vismed/MM/`

7 Conclusions and future work

The problem of occlusion when utilizing direct volume rendering in the scope of multimodal volume visualization of functional/anatomical data of the human brain has been overcome.

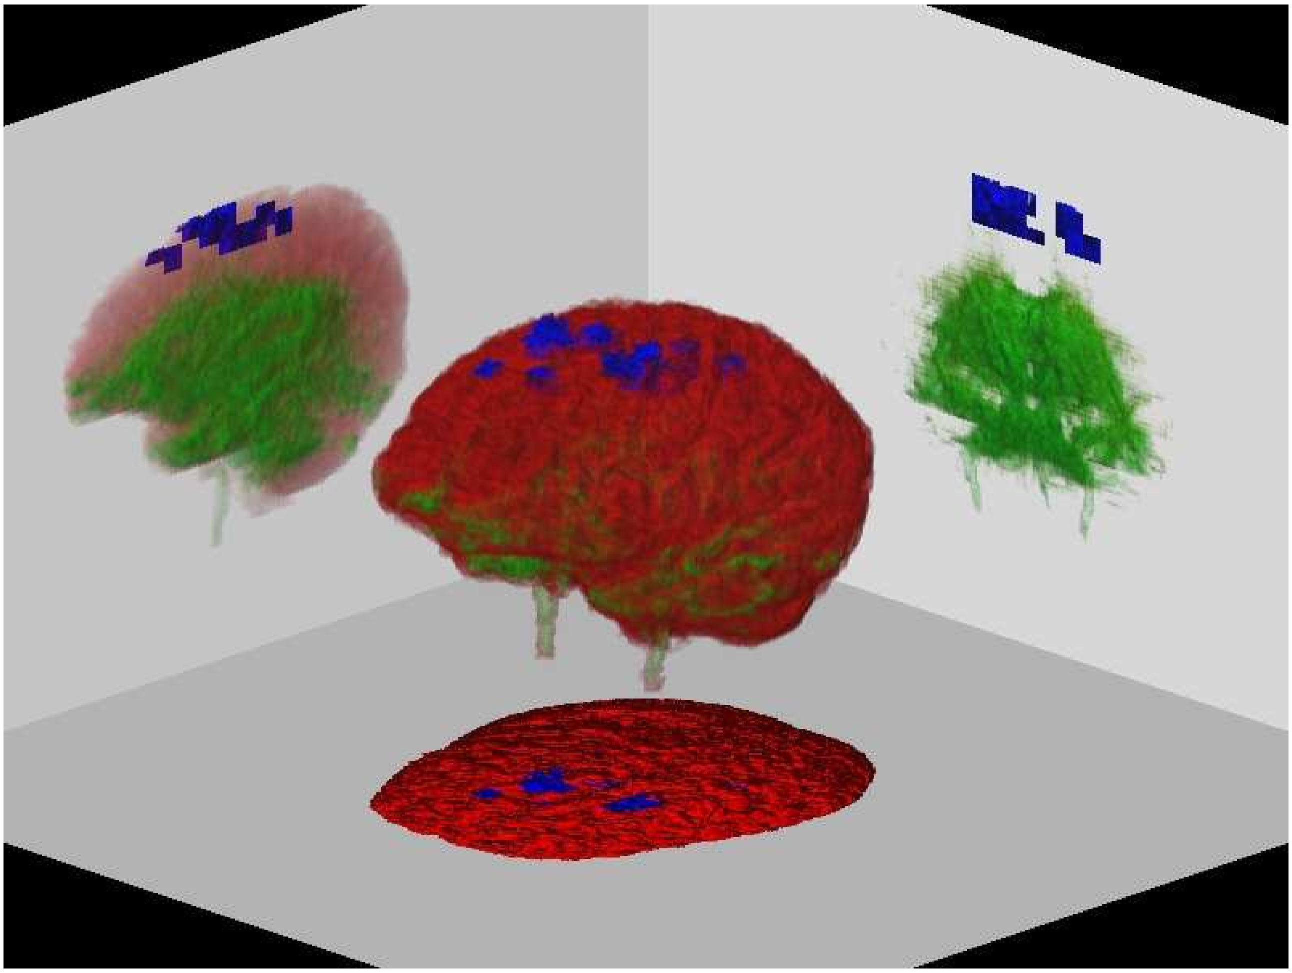

Figure 7: Magic Mirrors with different transfer functions (left and right mirrors) and projective view (bottom mirror)

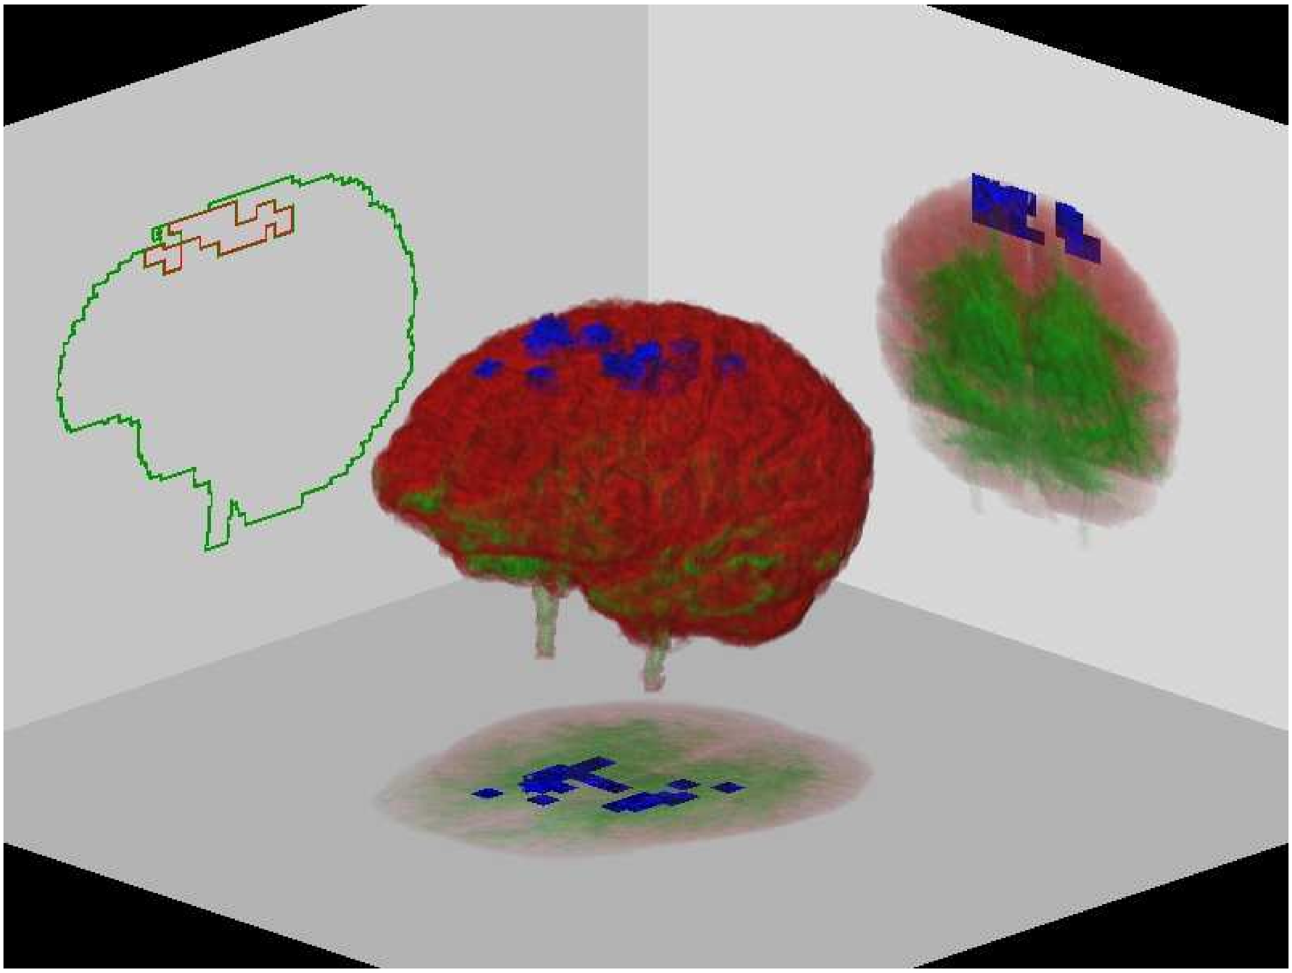

Figure 8: Magic Mirrors with contour display (left mirror)

Transfer-function volumes are used to realize highlighting and various volume sculpting tasks by assigning different transfer functions on a voxel by voxel basis. Multiple views of the visualized brain geometry in a single image provide additional information. Different representations of the multimodal scenario can be displayed by Magic Mirrors. Various combinations of data representations (utilizing different transfer functions and contour displays) provide the possibility to intuitively derive information like location and size of activated regions from the visualization.

Current shortcomings are addressed in our future research work. The possibility of interactive rotation of the visualization significantly aids the user in localizing the activated foci. A faster ray-casting system (shear warp based [9]) is going to be utilized, which should allow the interactive investigation and manipulation of the visualization. Currently the authors are working on the integration of the mirror planes into the shear warp approach. Regions of activation are going to be treated as entities, allowing to visually encode the magnitude of volume size and other properties in our direct volume rendering approach. Furthermore functional data sources obtained from different modalities are going to be included, in order to compare the results.

Finally, extensive acceptance tests with real world users are going to be performed. This should help to evaluate the practical value of the presented methods and to gain refined insights to the requirements in daily clinical practice.

8 Acknowledgements

The work presented in this publication has been funded by the V^is^M^ed^ project. V^is^M^ed^ is supported by *Tiani Medgraph*, Vienna, `http://www.tiani.com`, and the *Forschungsförderungsfonds für die gewerbliche Wirtschaft*, Austria, `http://www.telekom.at/fff/`. Please refer to `http://www.vismed.at` for further information on this project.

Additionally the authors would like to thank their project partners Prof. Dr. E. Moser and DI M. Barth, Vienna General Hospital, Institute for Medical Physics, for their invaluable contribution to the presented research work. Information on their work can be found via `http://lambda.meph.univie.ac.at`.

References

[1] R. Avila, T. He, L. Hong, A. Kaufman, H. Pfister, C. Silva, L. Sobierajski, and S. Wang. VolVis: A diversified volume visualization system. In R. Daniel Bergeron and Arie E. Kaufman, editors, *Proceedings of the Conference on Visualization*, pages 31–39, Los Alamitos, CA, USA, October 1994. IEEE Computer Society Press.

[2] D. W. Chakares and P. Schmalbrock. *Fundamentals of Magnetic Resonance Imaging*. Williams & Wilkins, 1992.

[3] K. J. Friston. Statistical Parametric Mapping. In *Functional Neuroimaging*, pages 79–93, Academic Press, San Diego, CA, 1994. R. W. Thatcher and M. Hallett and T. Zeffiro and E. R. John and M. Heurta.

[4] K. H. Höhne and A. Pommert. *Volume Visualization. Brain Mapping: The Methods*, chapter 17. A. W. Toga and J. C. Mazziotta, Academic Press, San Diego, CA, 1996.

[5] Victoria Interrante. Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In *Computer Graphics, Annual Conference Series*, pages 109–116, 1997.

[6] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Enhancing transparent skin surfaces with ridge and valley lines. In *IEEE Visualization '95*, pages 52–59, 1995.

[7] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *IEEE Visualization '96*. IEEE, October 1996. ISBN 0-89791-864-9.

[8] Victoria L. Interrante. Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 109–116. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.

[9] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. In *SIGGRAPH 94 Conference Proceedings*, Annual Conference Series, pages 451–458. ACM SIGGRAPH, July 1994.

[10] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, 8(3):29–37, May 1988.

[11] WE. Lorensen and HE. Cline. Marching cubes: a high resolution 3D surface reconstruction algorithm. *Computer Graphics (SIGGRAPH '87 Proceedings)*, 21:163–169, 1987.

[12] J. B. Antoine Maintz and Max. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.

[13] J. D. Mulder and J. J. Wijk. Parametrizable cameras for 3D computational steering. In W. Lefer and M. Grave, editors, *Visualization in Scientific Computing '97. Proceedings of the Eurographics Workshop in Boulogne-sur-Mer, France*, pages 165–176. Eurographics, Springer Verlag, Wien, New York, 1997.

[14] K. Rehm, K. Lakshminaryan, S. Frutiger, K. A. Schaper, D. W. Sumners, S. C. Strother, J. R. Anderson, and D. A. Rottenberg. A symbolic environment for visualizing activated foci in functional neuroimaging datasets. *Medical Image Analysis*, 2(3):215–226, 1998.

[15] K. Rehm, S. C. Strother, J. R. Anderson, K. A. Schaper, and D. A. Rottenberg. Display of merged multimodality brain images using interleaved pixels with independent color scales. *Journal of Nuclear Medicine*, 35:1815–1821, 1994.

[16] P. Rheingans. Opacity-modulating triangular textures for irregular surfaces. In *IEEE Visualization '96 Proceedings*, pages 219–225. IEEE Computer Society, October 1996.

[17] R. Sammouda, N. Niki, and H. Nishitani. Neural network based segmentation of magnetic resonance images. *IEEE Nuclear Science Symposium*, 4:1827–1831, 1995.

[18] J. Talairach and P. Tournoux. *Co-planar Stereotaxic Atlas of the Human Brain*. Thieme Medical Publishers, Inc., New York, 1988.

Multiple Views and Magic Mirrors - fMRI Visualization of the Human Brain

Andreas König, Helmut Doleisch, and Eduard Gröller

Vienna University of Technology
Institute of Computer Graphics
Karlsplatz 13/186/2, A-1040 Vienna, Austria
email:{koenig, doleisch, groeller}@cg.tuwien.ac.at

Category: research

Format: print

Contact:	Andreas H. Koenig
	Vienna University of Technology
	Institute of Computer Graphics
	Karlsplatz 13/186/2, A-1040 Vienna, Austria
phone:	+43 664 2647611
fax:	+43 (1) 587 49 32
email:	koenig@cg.tuwien.ac.at

Estimated # of pages: 8

Keywords: Volume Visualization, Multimodality, fMRI

Multimodal visualization of functional and anatomical data of the human brain is an important field in medical volume visualization. The aim of this application is to provide the user with information on the location of functional activations in the different regions of the brain. When the approach of direct volume rendering is chosen, the visual impression of details usually suffers from accumulating the image from colors and opacities derived from the data set. We present extensions to overcome this problem: *Transfer function volumes* are used for the highlighting of activated regions. *Multiple views* simplify the task of localizing these parts of the brain. Complementary information is brought into the visualization by *Magic Mirrors* in order to enhance the comprehensive view of the multimodal arrangement of volume data sets.