

Visualization of Time-Varying MRI Data for MS Lesion Analysis

Melanie Tory^{*}, Torsten Möller^{**}, and M. Stella Atkins^{***}

Department of Computer Science
Simon Fraser University, Burnaby, BC, Canada

ABSTRACT

Conventional methods to diagnose and follow treatment of Multiple Sclerosis require radiologists and technicians to compare current images with older images of a particular patient, on a slice-by-slice basis. Although there has been progress in creating 3D displays of medical images, little attempt has been made to design visual tools that emphasize change over time. We implemented several ideas that attempt to address this deficiency. In one approach, isosurfaces of segmented lesions at each time step were displayed either on the same image (each time step in a different color), or consecutively in an animation. In a second approach, voxel-wise differences between time steps were calculated and displayed statically using ray casting. Animation was used to show cumulative changes over time. Finally, in a method borrowed from computational fluid dynamics (CFD), glyphs (small arrow-like objects) were rendered with a surface model of the lesions to indicate changes at localized points.

Keywords: 4D visualization, multiple sclerosis, MRI, volume rendering, isosurface extraction, glyph

1. INTRODUCTION

Multiple sclerosis (MS) is a disease characterized by demyelination of neurons in the central nervous system. Scar tissue and demyelinated areas appear as bright lesions in spin-spin (T2) and proton density (PD) Magnetic Resonance Imaging (MRI) scans; thus, physicians use MRI images to diagnose MS and follow disease progression. One important diagnostic measure is lesion area (on a two-dimensional (2D) image) or lesion volume (in three dimensions (3D)). Current methods to determine these values involve time-consuming manual segmentation of MS lesions on each individual data slice. Furthermore, radiologists interested in disease progression must compare current images with older images of the same patient on a slice-by-slice basis. Since a typical diagnostic scan of 5mm thick slices results in approximately 22 slices each of two data sets (PD and T2), a full comparison of two time steps requires analysis of 88 2D images. Thus, radiologists currently do not have the time or the means to compare MRI scans from any more than two time steps (i.e. the current scan with the most recent previous scan), and cannot easily gain a long-term picture of disease progression. The objective of our study was to visualize changes in MS lesions over time. We explore various methods by combining ideas from volume rendering, isosurface rendering, and flow visualization.

2. RELATED WORK

1. Automated Segmentation

Recent work by members of our group (Atkins, Mackiewicz, and Krishnan) [2, 7] has produced methods for automating segmentation of the brain and MS lesions from PD and T2 MRI images. This eliminates the need for time-consuming manual segmentation, and thus makes possible fast, automated determination of lesion volume. Our method extends these approaches to calculate and display changes in lesion volume over time, and to render the segmented time series volume data in 3D.

2. Volume Visualization

Methods to visualize volume data have been extensively studied and developed. Isosurface extraction, such as that performed by the Marching Cubes algorithm, attempts to find surfaces within a volume and display them using traditional graphics methods [3, 9, 10]. An “isovalue” is arbitrarily chosen to represent the boundary between what is considered within the surface and what is considered outside the surface. Images generated using surface extraction methods provide information about overall structure in a 3D data set, but provide no information about what lies within the surface.

^{*} mktory@cs.sfu.ca; ^{**} torsten@cs.sfu.ca; ^{***} stella@cs.sfu.ca

By contrast, direct volume rendering attempts to visualize information throughout an entire 3D data set, so the user gains more information than is provided by a simple surface [3, 11]. Each point in the data set is classified according to which tissue it belongs, and is then assigned appropriate color and opacity values (opacity is a measure of how opaque an object is). A classified volume is rendered to the screen using one of many approaches (e.g. ray-casting or splatting). Current research in direct volume rendering focuses on improving both image quality and rendering speed. In addition, recent attempts have been made to find efficient data structures and algorithms for rendering time-series volume data [1, 4].

3. Flow Visualization

Although time-series data is sometimes considered in volume graphics, most research in this area comes from computational fluid dynamics, where various techniques have been developed to visualize flow of liquids and gasses. Glyphs are small pointed objects (e.g. arrows) that are placed within a scene and made to point in the direction of flow [11]. Glyphs can provide interesting information about flow direction at discrete points, but provide no information about flow between the points. To obtain detailed information, additional glyphs must be placed in the scene, which causes clutter. Other methods to visualize fluid flow include streamlines/timelines (lines consisting of a series of particles whose movement in the flow field are followed over time), texture advection (where small textures are warped in the direction of flow), and line integral convolution (an extension of texture advection).

A few recent attempts have been made to apply flow visualization ideas to tensor data from diffusion-weighted magnetic resonance imaging. Kindlmann and Weinstein use volume rendering with novel methods for assigning color and opacity values [6]. Laidlaw et al. take a different approach, using oriented "brush strokes" to visualize tensors, in a manner similar to glyph visualization or texture advection [8]. Finally, Zhang et al. [14] visualize tensors with the aid of stream tubes and stream surfaces.

3. METHODS

Three distinct visualization approaches (surface rendering, direct volume rendering, and glyphs) were implemented using the Visualization Toolkit (VTK)[11]. Each approach provided a unique view of the underlying data. Prior to visualization, however, significant data preparation was required. Details of preprocessing and visualization methods will be discussed in the following sections.

1. Data preparation

Time series data used in this study consisted of eleven time steps of 16 bit T2 and PD MRI brain scans of dimensions 256 X 256 X 22. Each slice of these data sets was segmented to remove the skull, scalp, and eyes from the brain tissue, using the technique described in [2]. Following this, the data was cast to 8 bit since our lesion segmentation program (see next paragraph) required data to be in this form. To ensure a good distribution of values in the 8-bit version, values three standard deviations below the mean intensity were mapped to zero and values three standard deviations above the mean were mapped to 255.

To simplify visualization of MS lesion changes, conspicuous and subtle lesions were segmented from each slice in the brain volume data sets using the method described in [7]. Output of the segmentation program consisted of two binary lesion masks, one for conspicuous lesions and one for all lesions (conspicuous plus subtle). To convert this data to a non-binary grey-scale, any voxels with value 1 were replaced by the value at the same position in the eight bit T2 or PD brain data. Essentially the segmentation output was used as a mask to select voxels from the T2 and PD MRI data sets.

Since each data set was acquired in a separate session, position of the subject will not correspond exactly between time steps; hence, registration is necessary. Registration is particularly important since MS lesions are relatively small relative to the brain as a whole. A slight change in position of the head could severely affect analysis of lesions within each data slice, since misregistration artefacts could be as great or greater than actual lesion changes [5].

In all cases, registration was performed using the Automated Image Registration (AIR) software package, version 3.08 [12, 13]. Registration parameters were derived using the AIR program "alignlinear" with the segmented brain volumes as input (not the segmented lesions). A rigid body six-parameter model was used and noninteraction of spatial parameter derivatives was assumed (refer to the AIR documentation for an explanation of AIR arguments [12]). Threshold values of 10 were selected for all data sets, and all other AIR default values were used. Following this, the derived parameters for each time step and volume type (PD or T2) were applied to the brain data and to both lesion data sets for that time step and volume type, using the AIR programs "mv_air" (to select which file to register) and "reslice" (to perform registration). All default

AIR argument values were used here, except that the data sets were not interpolated to cubic voxels; instead, voxel size was kept the same as the input volumes (256 X 256 X 22) so that volumes that were registered would be the same size as the volume they were registered to. The resulting registered data was used as input to our visualization programs.

2. Surface rendering

In one approach, we chose to display lesions as 3D contours using traditional computer graphics techniques. Isosurfaces of MS lesions were extracted using the marching cubes algorithm. Isovalues were selected based on histograms such that all lesion tissue in the segmented data sets would appear. Lesion data from one or more time steps could be displayed in the same scene; each time step was displayed in a different color to emphasize changes. In addition, an isosurface of the entire brain (extracted from the registered but unsegmented T2 data) was displayed to help orient the lesions in terms of their spatial position. All isosurfaces were given opacity 0.3 so that they would be translucent, allowing the user to see other isosurfaces below the uppermost one. User interaction was implemented such that a user could rotate, zoom, and pan the 3D display, as well as save an image in any orientation.

A variation of the above program allowed a user to select a single viewing orientation and produce a series of images for off-line viewing as an animation showing changes over time (from one data step to the next). Because the number of time steps in our data was small, we used linear interpolation between real data sets to produce more frames for a smoother animation. Furthermore, because the program was designed to produce a set of images for off-line viewing, user interaction was not enabled.

3. Direct volume rendering

A second approach made use of direct volume rendering, in which all points in the volume contribute to the final image so that more than a simple surface can be seen. However, instead of rendering the lesions themselves, voxel-wise differences between two time steps were calculated and the resulting intensity change was rendered. A color and opacity map was designed to indicate magnitude and direction (positive or negative) of the change at each voxel. Note that because we segmented the lesions in each time step, non-lesion material has an intensity value of zero; hence, size changes appear as very large intensity changes. User interaction and a surface model of the entire brain were incorporated into the program similar to the surface model program described above.

As with surface rendering, a variation of the direct volume-rendering program was designed to produce images for off-line animation of cumulative lesions changes. Linear interpolation between time steps was again used to increase the number of frames, and user interaction was not enabled.

4. Glyphs

In the final approach, small glyphs were used to illustrate the direction (increase or decrease) and strength of MS lesion changes. First, changes in lesion data between two time steps were found by subtraction. An isosurface was extracted from this change data using the marching cubes algorithm, and displayed for orientation purposes. This surface can be thought of as an "isosurface of change", and was intended to point out areas where change (either positive or negative) was higher than a threshold specified by the user. Glyphs were displayed on the same image to show changes at localized points, both on the isosurface and elsewhere.

In one model, glyphs were designed to point in one direction if there was a positive change at the point or the exact opposite direction if the change was negative. They were scaled according to the magnitude of the change, such that a small change resulted in a small glyph and a large change in a large one. In addition, glyphs were colored according to whether the change was negative or positive. We believed that redundant mapping of glyph color and direction to change direction (negative or positive) would emphasize this feature in the data. A second glyph model was similar, but had glyphs point in the direction of the surface normal instead of simply right or left, and did not scale the glyphs according to the amount of change at the point. The idea here was that lesion areas might be thought of as "growing outwards" in the direction of the surface normal or "shrinking inwards" in the opposite direction. Note, however, that we do not know for sure whether this is really the case since we have not consulted with radiologists on this issue.

In both glyph programs, a surface model of the entire brain could optionally be displayed along with the MS lesions and glyphs for orientation purposes. In addition, the programs were interactive so that the brain and glyphs could be rotated, zoomed, and panned.

4. RESULTS AND DISCUSSION


1. Summary Statistics

Summary statistics were computed on each data set and compiled into Table 1. Volume values are the sum of MS lesion areas for each slice in the dataset multiplied by slice thickness, and mean values are the average of all non-zero intensities in the volume (i.e. the mean intensity of the lesions, not including the background).

Table 1: Summary statistics for time series data (average intensity values and volume in cm^3)

Time Step	Conspicuous Lesions			All Lesions		
	Vol (cm^3)	T2 Mean	PD Mean	Vol (cm^3)	T2 Mean	PD Mean
1	40.1	194	195	66.5	177	184
2	37.2	197	194	64.2	178	185
3	37.6	196	194	65.8	176	186
4	35.8	199	194	61.1	178	186
5	35.2	198	195	61.6	176	186
6	34.8	195	193	60.3	178	185
7	35.6	196	192	62.0	176	184
8	34.1	196	193	59.9	175	185
9	32.7	194	192	56.4	176	184
10	32.7	196	192	57.1	176	183
11	32.0	196	193	56.3	178	185
Mean	35.3	196	193	61.0	178	185

Figure 1: Changes in total MS lesion volume



As shown in Figure 1 and Table 1, total lesion volume decreased over time, by about 20% for conspicuous lesions and 15% for all lesions. Meanwhile, average PD and T2 intensities of the remaining lesion areas did not change. However, these summary values give no information about localized changes. For example, lesions in one part of the brain could expand while others shrink, and areas of increased intensity could balance areas of decreased intensity. We used 3D visualization techniques to explore such spatial variations.

2. Surface rendering

Figure 2 illustrates images generated by the surface extraction algorithm. (Figures 2-5 also appear on a color plate in the paper proceedings.) Each lesion data set is a different color, such that lesions from the first time step are blue, lesions from the last time step are red, and middle time steps are assigned varying shades of purple. Areas containing lesions from more than one time step appear as color combinations. These displays provided a 3D impression of how lesions were oriented within the brain, the amount of volume they occupied, and where large and small lesions were located. Since any number of data sets could be displayed at once, it was possible to see that some older lesions had disappeared (blue areas in Figure 2) and some new lesions had appeared (red areas in Figure 2 (a)).

Note however, that because the 3D objects are projected into two dimensions, color combinations may be caused by lesions lying behind one another rather than in exactly the same spot. Rotation of objects on the screen allowed the user to gain additional understanding of 3D orientation, contributing to the resolution of this problem. It was also noticed that with all eleven time steps displayed (see Figure 2 (c)), the image appeared almost entirely red, giving the impression that a large number of new lesions had appeared; however, this coloring was actually caused by the fact that the opacity used (0.3) was too high. Thus the last four data sets in the time series (all in shades of red) occluded all previous data sets (in shades of blue and purple). Lowering the opacity solves this problem, but makes it hard to see areas with only a single lesion. Thus, we believe that interactive tools that allow the user to assign each time step its own opacity and isovalue could significantly improve the usability of the system.




Figure 2: Surface rendering results for T2 conspicuous lesions. The marching cubes algorithm was used to extract isosurfaces from several registered lesion data sets. All isosurfaces were then rendered to the same window using VTK. An isosurface of the entire brain was rendered in the same way, for orientation purposes.

Animations showing changes in lesion surfaces allowed all eleven time steps to be viewed without occlusion; however, since the animation program was not interactive, the brain could not be rotated or zoomed and the program had to be re-run whenever the user wanted to change parameters. An interactive animation is the obvious solution to this problem.

3. Direct volume rendering

Surface models provided information about lesion volume and appearance and disappearance of individual lesions, but no information about localized intensity changes. This information may be useful in diagnosis, although the exact meaning of intensity changes in white matter is not yet known. Like the manual methods, our lesion segmentation method relies on the high intensity of the lesion pixels in the PD and T2 images; spectroscopic evaluation of MS lesions and normal-appearing white matter is a current research area. To attempt to visualize intensity information, we used direct volume rendering.

Direct volume rendered images of MS lesion changes are displayed in Figure 3. Areas of strong increase appear red, areas of strong decrease appear blue, and relatively unchanged lesion areas appear green. Slightly changed areas appear as linear combinations of appropriate colors. Like the surface model program, the direct volume rendered image contains a surface model of the head, and users can rotate objects in the scene to improve spatial orientation. Animation can be used to display cumulative changes as increasingly bright reds and blues, but like the surface animation, user interaction is lost.




Figure 3: Direct volume rendered images comparing T2 intensity of time steps 2 and 4 of (a, c) conspicuous lesions and (b) all lesions. Areas of increased intensity are shown in red, decreased intensity in blue, and relatively unchanged intensity in green. (c) Side view appears blurry because of low resolution in one dimension of the data.

Although we found that overall intensity of the lesions did not change (see Table 1), Figure 3 illustrates that localized changes did occur. For example, in Figure 3 (a), the blue spot in the middle represents an intensity drop, while the red and brown spots to the left and right represent intensity increases. Notice that the blue spot has disappeared in the view containing all lesions (Figure 3 (b)). In this case, subtle lesions not contained in the conspicuous lesion data sets occluded the intensity drop within the conspicuous lesions. This became obvious on rotation of the image when the change in viewpoint reduced the level of occlusion such that the blue spot could once again be seen. Lowering the opacity of each data point may have also prevented occlusion; thus, we believe that a method to interactively change transfer functions (functions that assign opacity and color values to voxels) would be very useful for data analysis.

In addition, we found the poor resolution in one dimension of our data sets to be a problem. Resolution was 256 X 256 within each slice, but there were only 22 slices in each data set. The first obvious problem here is that some lesions could be missed altogether in an MRI scan, such that they could seem to appear or disappear from one scan to the next. Furthermore, Figure 3 (c) illustrates that volume rendering these data sets with the viewpoint at the side results in very blurry images that are difficult to understand.

4. Glyphs

Figure 4 and Figure 5 show images of glyph models 1 and 2, respectively. In both cases, glyphs appear red if the change at that point was positive and blue if the change was negative. In Figure 4, glyphs are scaled according to the strength of the change and point in opposite directions depending on whether the change was positive or negative. Glyph direction and size alone provide all the information, but it was believed that the addition of a color dimension would increase the speed at which information was conveyed.




Figure 4: Glyphs displaying MS lesion intensity changes between time steps 2 and 10, shown with an isosurface of the entire brain. Glyphs at points of increased T2 intensity are red and point right, and glyphs at points of decreased T2 intensity are blue and point left. Length of glyph indicates magnitude of change. Inset shows part of the image after a zoom operation.

Figure 5 is similar, but glyphs point in the direction of the surface normal and are not scaled according to strength of the change. The hypothesis here was that glyphs might grow and shrink in the direction of the surface normal; of course, this may or may not be the case in reality. In both glyph models, an isosurface of change indicated areas where change magnitude was greater than a user-specified threshold. In addition, users had control over what factor to scale glyphs by (how big glyphs should be relative to the brain isosurface) and what fraction of glyphs to display. For example, in Figure 5 (a) a low zoom factor is used, so to avoid excessive clutter and ensure glyphs are visible, one would choose a higher scaling factor and a smaller fraction of glyphs to display than for Figure 5 (b), which has a much higher zoom factor. A problem arose, however, since these parameters were specified ahead of time and could not be changed while the program was running. As a result, as users zoom in and out, scaling of the glyphs (relative to the brain) becomes inappropriate. For

example, the inset in Figure 4 shows glyphs that appear excessively large when the image is zoomed. Information could be conveyed much better if glyph fractions and scaling either (1) changed automatically as the image was zoomed in and out, or (2) could be interactively changed by the user during program execution.




Figure 5: Glyphs displaying MS lesion intensity changes between time steps 2 and 10. Glyphs at points of increased T2 intensity are red, and those at points of decreased T2 intensity are blue. Glyphs point in the direction of the surface normal. (a) Zoom factor 1, scaling 0.01, 1/25 of all glyphs are displayed. An isosurface of the entire brain is also shown. (b) Zoom factor 15, scaling 0.002, 1/6 of all glyphs are displayed.

Informal evaluation of our results by two physicians revealed that they were more comfortable with surface and volume rendered images than with our glyph images. They felt that the glyphs conveyed a false impression of spatial movement, and were unlike any images they were accustomed to viewing. Training and improved user interaction may or may not alleviate such problems in the future.

5. CONCLUSIONS AND FUTURE WORK

Visualization of changes in MS lesions is an interesting problem that is yet to be fully explored. This study attempted to evaluate several visualization methods in terms of their application to three-dimensional MS lesion visualization. Surface models provided useful information about spatial orientation of lesions and changes in lesion volume, but provided no information about intensity changes. Direct volume rendered images of lesion changes addressed this problem by illustrating spatially localized changes of intensity. Glyph images were used to approach the problem from a different perspective. Glyphs clearly illustrate lesion changes at localized points in the brain, but come at the cost of a very cluttered view. Further work is needed to optimize these images to convey information effectively. In addition, we believe that the development of graphic interfaces and improved user interaction will improve the effectiveness of our visualizations. Other future work includes optimizing the programs for speed, exploring other visualization techniques, re-examining our arbitrarily chosen color scales to see if they can be improved, looking at other imaging modalities and applications, and iterative development and analysis with users.

ACKNOWLEDGMENTS

The authors wish to thank Craig Jones for providing time series data. The assistance of Wilf Rosenbaum and Kevin Siu with existing segmentation programs is much appreciated.

REFERENCES

1. Anagnostou, Kostas, Tim J. Atherton, and Andrew E. Waterfall, "4D Volume Rendering With The Shear Warp Factorisation," *Volume Visualization and Graphics Symposium 2000*, pp. 129-137, Oct. 2000.

2. Atkins, M.S. and Blair T. Mackiewich, "Fully Automatic Segmentation of the Brain in MRI," *IEEE Transactions on Medical Imaging*. 17(1), pp. 98-107, Feb. 1998.
3. Bankman, Isaac N., ed., *Handbook of Medical Imaging: Processing and Analysis*, Academic Press: San Diego, California, 2000.
4. Ellsworth, David, Ling-Jen Chiang, and Han-Wei Shen, "Accelerating Time-Varying Hardware Volume Rendering Using TSP Trees and Color-Based Error Metrics," *Volume Visualization and Graphics Symposium 2000*, pp. 119-128, 155, Oct. 2000.
5. Hajnal, Saeed, Oatridge, Williams, Young, and Bydder, "Detection of Subtle Brain Changes Using Subvoxel Registration and Subtraction of Serial MR Images," *Journal of Computer Assisted Tomography*. 19(5), pp. 677-691, 1995.
6. Kindlmann, Gordon, and David Weinstein, "Hue-balls and Lit-Tensors for Direct Volume Rendering of Diffusion Tensor Fields," *IEEE Visualization 99*, Oct. 1999.
7. Krishnan, K. and M.S. Atkins, "Segmentation of Multiple Sclerosis Lesions in MRI – An Image Analysis Approach," *Proceedings of SPIE International Symposium on Medical Imaging*, Feb. 1998.
8. Laidlaw, David H., Eric T. Ahrens, David Kremers, Matthew J. Avalos, Russell E. Jacobs, and Carol Readhead, "Visualizing Diffusion Tensor Images of the Mouse Spinal Cord," *IEEE Visualization 98*, pp. 127-134, Oct. 1998.
9. Levoy, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*. 8(5), pp. 29-37, May 1988.
10. Lorensen and Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," *Computer Graphics*. 21(3), pp. 163-169, July 1987.
11. Schroder, Martin, and Lorensen, *The Visualization Toolkit*, 2nd ed. Prentice Hall PTR: New Jersey, 1998.
12. Woods, *Automated Image Registration Documentation*. <http://bishopw.loni.ucla.edu/AIR3/index.html>.
13. Woods, Cherry, and Mazziotta, "Rapid automated algorithm for aligning and reslicing PET images," *Journal of Computer Assisted Tomography*, 16, pp. 620-633, 1992.
14. Zhang, Song, Charles T. Currey, Daniel S. Morris, and David H. Laidlaw, "Streamtubes and Streamsurfaces for Visualizing Diffusion Tensor MRI Volume Images," *IEEE Visualization 2000: Work in Progress*. Oct. 2000.

COLOR PLATE




Figure 2: Surface rendering results for conspicuous lesions. The marching cubes algorithm was used to extract isosurfaces from several registered lesion data sets. All isosurfaces were then rendered to the same window using VTK. An isosurface of the entire brain was rendered in the same way, for orientation purposes.




Figure 3: Direct volume rendered images comparing T2 intensity of time steps 2 and 4 of (a, c) conspicuous lesions and (b) all lesions. Areas of increased T2 intensity are shown in red, decreased T2 intensity in blue, and relatively unchanged intensity in green.




Figure 4: Glyphs displaying MS lesion changes between time steps 2 and 10, shown with an isosurface of the entire brain. Glyphs: red and point right = increased intensity, blue and point left = decreased intensity. Length of glyph indicates magnitude of change. Inset shows part of the image after a zoom operation.




Figure 5: Glyphs displaying MS lesion changes between time steps 2 and 10. Glyphs at points of increased intensity are red, and those at points of decreased intensity are blue. Glyphs point in the direction of the surface normal. (a) Zoom factor 1, scaling 0.01, 1/25 of all glyphs are displayed. An isosurface of the entire brain is also shown. (b) Zoom factor 15, scaling 0.002, 1/6 of all glyphs are displayed.