

Angiotool: A tool for interactive visualization of MRI vector and tensor fields

Abhir Bhalerao
Department of Computer Science
University of Warwick
Coventry
UK
abhir@dcs.warwick.ac.uk

Paul Summers
Clinical Neurosciences
Kings College Medical School
London
UK
p.summers@iop.kcl.ac.uk

May 30, 2001

Abstract

MR imaging is a uniquely powerful tool for non-invasively mapping blood flow and molecular diffusion. These properties have clinical significance but the lack of efficient and intuitive, display and exploration tools for the clinical expert is an obstacle to bringing these images into diagnostic use. We have developed a visualization application (Angiotool) for vector and tensor field data derived from MR phase contrast angiographic and diffusion tensor image (DTI) sequences which augments routine display functions with both quantitative and qualitative exploration features. One of its unique capabilities is to perform interactive tracking of streamlines in both velocity and diffusion tensor data. We describe the GUI model used by Angiotool and illustrate its interactive capability with example results from clinical cases. Central to the design of Angiotool is the ability to interact and navigate through the data in both 2D and 3D displays. Although implementations may differ, we see our interaction model as well suited to new applications of this type.

1 Introduction

Motion, whether due to blood flow or molecular diffusion can have measurable effects on both the phase and cumulative magnetisation arising from hydrogen nuclei (spins) in the body [1]. Developing ways and means for visualising and analysing these data that meet the needs of radiologists and surgeons has been a goal of researchers following various paths, including: visualising and tracking blood flow data [2, 3]; vascular segmentation [4, 5, 6]; and analysing diffusion data [7, 8, 9, 10].

Central to MR imaging is the dependence of a spin's phase on the time integral of the dot product between its position and the vector describing the local magnetic field gradient (or simply gradient). Adequate control of the imaging gradients allows the vector-valued magnetization (typically represented in complex value form) to be mapped by use of simple Fourier transforms of the detected signals. Independent of the mapping process, applying an additional bipolar gradient consisting of two successive gradient applications of equal duration and strength, but with opposite polarity, allows the phase of the spins to be related to its position. A stationary spin will yield no overall phase shift. A spin which changes position over the course of the bipolar gradient application however, will undergo a phase shift proportional to its displacement and the timing of its movement. When the spins within a voxel move at constant velocity throughout a bipolar gradient their resulting phase shifts are linearly related

to the component of their velocity along the gradient direction. In the presence of a strong gradient, highly random motion of the spins, i.e. molecular diffusion, can be imaged [7]. As velocity is a vector quantity, three component images (plus a reference unencoded image) are needed to produce a velocity field. Similarly, diffusion in a locally homogeneous material is described by a 2nd order tensor (a 3x3 symmetric matrix) and requires at least seven images for its description.

The velocity or diffusion component maps may be further processed to produce summary images. A common endpoint of phase contrast angiography (PCA) is an angiographic speed image produced by taking a square root of the sum of squared the velocity components. The commonest way of post-processing diffusion images is to determine the net or anisotropic diffusion coefficient (ADC) by using the principal eigenvalue of the diffusion tensor. In general, these images provide a concise description of some aspect of the vector or tensor field. As well as the calculated and derived images, it is typical for an anatomical (modulus image from the unsensitized data acquisition) to be reconstructed also. Altogether, a typical PCA sequence will produce of 5 image sets of size 256x256x120 voxels while a diffusion dataset might consist of 7 image sets of size 128x128x60 voxels. The vastness of the data sets pose unique computational challenges.

In this article, we outline clinical uses of MR motion imaging which motivate our work toward an intuitive graphical tool that meets the requirements of clinical experts. The GUI design model of our application, Angiotool, is then described. Its display and interactive analysis functions, which can be used equally for vector valued velocity and tensor data, are illustrated from clinical cases which highlight of Angiotool's capabilities.

2 Use of images and Existing tools

Clinically, the attraction of PCA has been its ability to depict flowing blood without contamination from static tissues. As distinct from other methods of MR angiography, PCA also has the ability to delineate flow patterns and quantify velocity. Diffusion imaging on the other hand is most commonly associated with the depiction of strokes, where restricted diffusion indicates recent occurrence (acute stroke and uncertain fate of the tissue) and elevated diffusion is seen when dead neurons are replaced by cerebrospinal fluid (chronic stroke). Considerable further interest in diffusion imaging is associated with the potential to identify patterns or pathways of connectivity within the brain on the basis of how diffusion anisotropy reflects the course of myelinated (message conducting) neurons. There are similarities in the specific questions asked by experts for the two types of MR images:

- is a vessel patent / is a neuronal tract intact?
- what is the degree of stenosis (narrowing) / is an ischaemic (stroke) lesion old or new?
- does a particular vessel feed or drain a given region / what are the terminal connections of a neuronal tract?
- what is the flow pattern in a given region / what is the connectivity of a cortical region?
- how has a flow pattern been affected post-operatively / is the neuronal tract influenced by adjacent pathology?

The major medical equipment vendors and several independent imaging software developers provide medical image viewing and analysis packages. Typically, the underlying software is bundled with a hardware system. This in part relates to regulatory practice, and in part

to practicalities of customer support and maintenance. This approach is exemplified by Vitrera(Voxelview) (Vital Images Inc. Minneapolis Minn.) which takes full advantage of hardware and software acceleration to facilitate rapid view rendering. Perhaps the most widely-used, multi-platform software-only approach is Analyze (Analyze, Mayo Clinic, USA). Whereas Voxelview is designed particularly for volume rendering and visualization, Analyze incorporates a wide range of additional tools for such tasks as segmentation, image registration, format conversion, surface rendering, overlays of two image sets, and the measurement of distances, angles and areas. Both packages support the loading of images from a number of scanner types, allow the generation of orthogonal and obliquely reformatted images, and provide maximum intensity projection (MIP) and other rendering tools. A notable difference in approach is seen in the generation of dynamic “fly-throughs” of rendered displays. Whereas Analyze has relied on scripting tools to control object and viewer pose, Voxelview uses a point and click interactive approach which records the viewer’s movements through and around the dataset. We feel the latter approach better represents the type of intuitive interaction we hope to achieve within Angiotool. Neither Analyze nor Voxelview support the viewing or manipulation of vector or tensor datasets.

As yet, there is no universally accepted paradigm for user interaction with and visualisation of vector and tensor data for clinical use. In fact, to our knowledge, none of the commercially available platforms for medical image viewing deal with vector or tensor data as such. In general, do they allow interaction with more than a single 3D image volume at a time except for use in image registration or multispectral segmentation routines. In fact, of the 5 or more datasets associated with a PCA study, only the derived “speed” images receive attention - typically through use of a MIP display. For tensor data, ADC and fractional anisotropy may be rendered in similar fashion(e.g. [11, 9]). In neither case does the MIP of these derived scalar metrics convey the directional information contained in the acquired data. Moreover, MIP viewing is performed independently of display of the corresponding tissue images, which complicates the task of determining relationships between vascular and non-vascular structures. A final restriction of most clinical image viewing packages, is that quantitative analysis of the vector or tensor data is often precluded by discarding the underlying data for the brevity of the derived scalar images.

3 Angiotool: Data Visualization and Navigation

We have attempted to develop a GUI based approach to interactive visualization of vector and tensor data which meets clinical needs by incorporating both traditional cut plane and MIP displays of derived images, with the added abilities to access the underlying data through quantitative and dynamic qualitative visualization. The GUI presented by Angiotool can be decomposed into three main displays areas: a square 3D rendering window, an adjacent 3D analysis window of equal size, and a triplet of smaller windows displaying slice-by-slice views of the data (Figure 1). Quantitative information is displayed in a scrolling text window and, where appropriate, auxiliary graphical windows.

3.1 Orthogonal Slice Views: Tissue and MIP images

Most radiologists prefer the light-box and film paradigm of slice-by-slice viewing to sole reliance on rendered displays of 3D data. Viewing in this manner allows better (if laborious) determination of spatial relationships between structures which may be obscured in static 3D views, and is widely used to confirm impressions even where for instance MIPs are used to gain an overview of the data. The generation of orthogonal planar images from 3D data, and

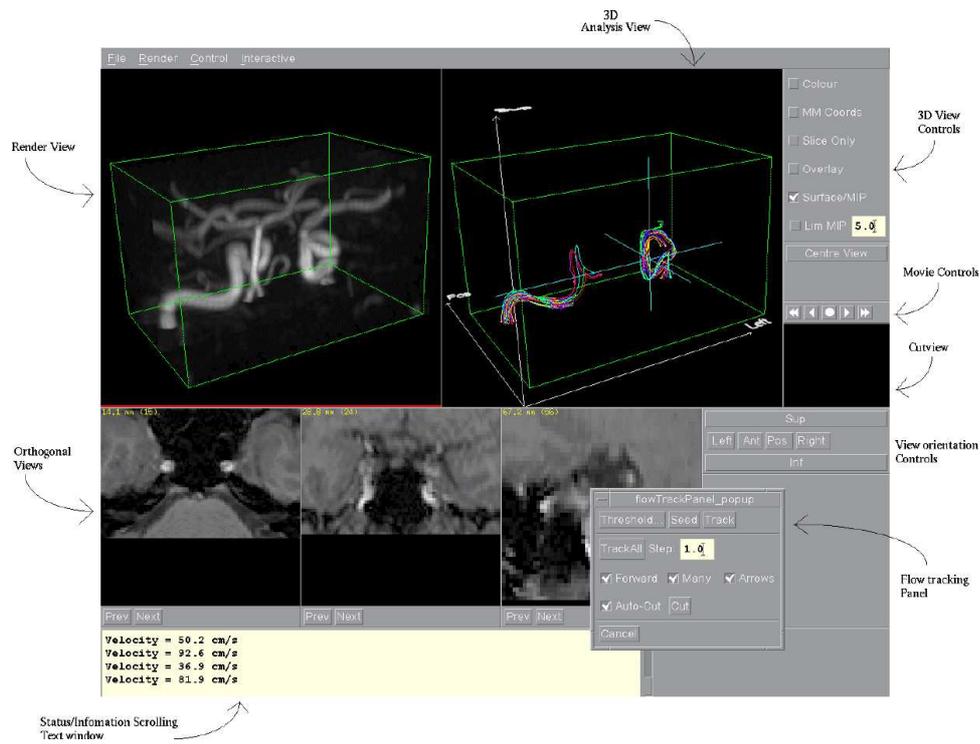


Figure 1: Overview of GUI layout of Angiotool. The user is presented with two square 3D rendering windows (top) and a triplet of smaller orthogonal views (centre) of the data. The left-hand 3D windows display an oblique MIP of the speed of a cerebral PCA. The right hand window shows the results of a flow tracking experiment in the right-internal carotid artery. The 3D cursor can be moved by selecting points in the MIP image (which looks up the corresponding 3D coordinate using a depth buffer) e.g. select the seed point for flow tracking.

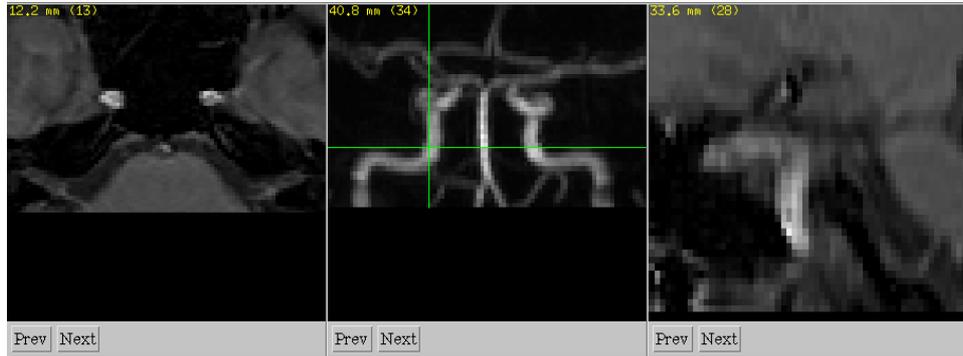


Figure 2: Orthogonal slice viewers can alternatively show MIPs in respective directions. Anatomical data is being viewed in the left-most and right-most displays; functional data in the centre display. The point-and-click 3D cursor navigation continues to operate when MIP view is shown allowing users to manually track along a vessel or fibre tract to display orthogonal cross sections at these points in associated tissue volume. Next/Previous buttons enable explorations of slices either side of current cursor position.

their simultaneous display is a powerful yet comfortable extension of the traditional clinical approach which is facilitated by computer display. In Angiotool, as with other orthogonal slice viewers, the views are linked by a 3D cursor (Figure 2). This allows for a simple point and click navigation through the data: a click in any of the three views will update the other two with the ‘cross-section’ associated with X and Y co-ordinates centred on the click position. Such a display model is widely used in CAD and architecture where centre lines link the two elevation and plan views. Within Angiotool, we typically use the T1 weighted tissue images from PCA data or either the T2 weighted tissue or derived ADC image from diffusion studies to provide comprehensive overviews of general anatomy in the orthogonal slice viewer. In fact, any co-registered data set, such as CT, can be nominated as the ‘anatomical’ image and viewed simultaneously with the phase contrast (functional) image. For each of the orthogonal view directions, a corresponding MIP of the speed (PCA) or ADC (diffusion) data is precomputed. With a single mouse click the user can toggle between the MIP and anatomical slice data. To aid relating the anatomical view to the angiographic or diffusion structures of interest, the point-and-click cursor navigation continues to operate on the MIP image. With this, when an image feature such as a vessel, is manually tracked with the cursor in the MIP view, the other two views are automatically re-centred on the cursor position. Within the orthogonal slice viewer, view-independent next-slice and previous-slice buttons allow the user to make limited explorations on either side of an initial location.

3.2 3D display

The two square, 3D display windows show the same viewpoint with the orientation controls and cursors being linked together. The left hand window (**3D render window**) shows a ray-cast representation of the data volume either as a surface shaded display (gradient shaded) or as a perspective MIP as shown in Figures 3(a) and (c). Full brute-force ray casting is computationally expensive, so the render window performs an incremental ray casting in a multi-resolution fashion. Rendering accuracy is temporarily traded-off against speed to allow the user to quickly re-orient themselves to a desired viewpoint and image zoom setting, achieving the low latency, continuous feedback interaction essential for any effective 3D navigation control [12]. The multi-threaded implementation seamlessly scales over multi-processor SMP architectures [13].

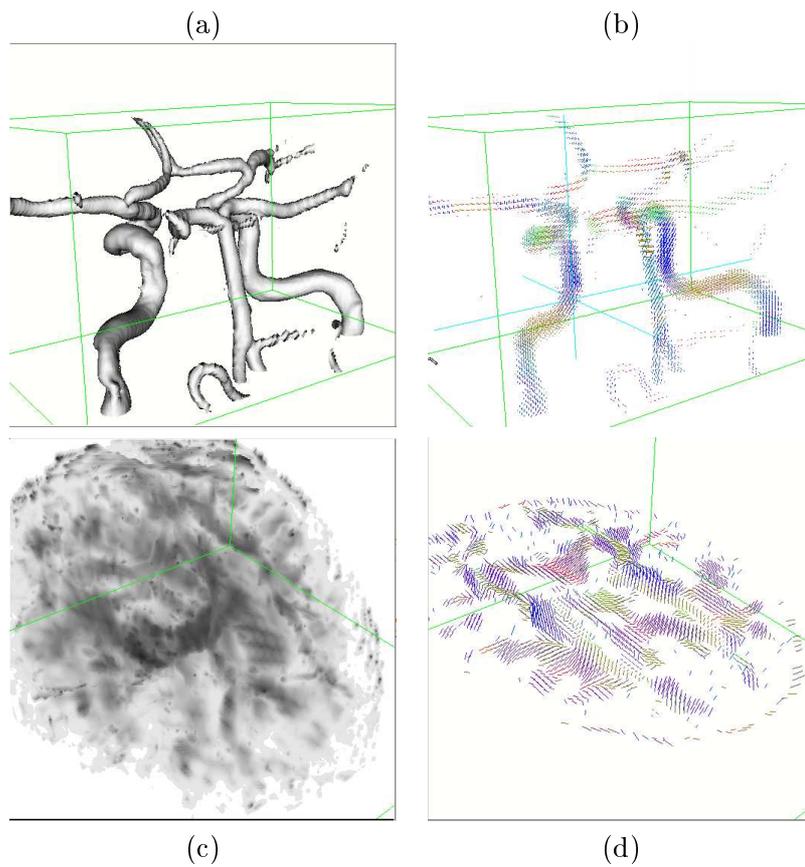


Figure 3: (a) Surface rendering of a speed iso-surface depicting major vessels from cerebral MRA. (b) Velocity data of (a) shown as a field of vector where the colour represents vector orientation. (c) MIP of ADC from an example diffusion image. (d) Plot of principal eigen vectors across a slice from diffusion image (c). Neuronal track divergence, cross-hemisphere and anterior-posterior connectivity is apparent.

The render window maintains a depth buffer (or z buffer) of the current projection from which the (x, y, z) cursor position can be mapped. When each MIP value (or surface colour) is projected onto the view plane, its voxel position is recorded in a 2D buffer $Z(i, j) = (x, y, z)$. When the user selects a point $[i, j]$ in the MIP using the cursor, Angiotool ‘looks-up’ the appropriate voxel address (x, y, z) . This may be used to redefine the slice positions in the orthogonal viewer or to provide initial points for analysis as described below. Because the projected value (x, y, z) in a MIP may be from brighter values either in front of or behind the vessel of interest. To prevent such depth ambiguity errors when tracking a vessel, a limited-MIP where only voxels which lie between two planes z_{min} and z_{max} parallel to the view plane are projected can be used. By reducing this depth value to be about the size of the voxel dimensionality (e.g. 1mm), the projected image reduces to a cut slice through the current cursor position providing a means of generating oblique slices.

The right hand window (**3D analysis window**) is used to display results of analyses on the vector or tensor data volume. Whereas the render window displays voxel information by ray casting, the analysis window displays graphics geometry (points, lines, planes and text). The 3D cursor, the bounding navigation cube, axes and data orientation label objects are ever present in analysis window. The natural association of vectors with line segments or arrows provides a simple means of conveying the underlying data for PCA. The flow field may be displayed by rendering the 3D vector field as a grid of small lines through the centres of each voxel. The vector direction and magnitude are colour coded as follows:

- Velocity orientation is coded such that each component (v_x, v_y, v_z) is mapped to a corresponding colour channel (r, g, b) e.g. $r = (v_x + \max(|v_x|))/2$, normalised to the appropriate colour channel range. This mapping has the effect that all vectors with the same orientation have the same colour.
- Velocity magnitude (i.e. speed of flow) proportionally controls the length of the line representing the vector.

A user defined threshold is used to avoid clutter from stationary tissues and air permitting vector display as illustrated in Figure 3(a). Several authors have used ellipsoids as a voxel-wise display analogy for diffusion tensor data (e.g. [8]). As much of the attention in our institutions has focussed on questions of connectivity, we have simplified the ellipsoid model to one of displaying the principle eigenvector of the diffusion tensor for each voxel 3(d).

Direct rendering of vector data at the voxel level is a first step to making use of the available information. This approach is however, not without its limitations. For instance, the identification of likely paths through the data - blood flow streamlines or neural connectivity, must be inferred. Also, structural boundaries are often difficult to distinguish in the vector display than in the anatomical or MIP images. The anatomical rendering from the render window can be superposed onto contents of the 3D analysis window. This allows the user to visually fuse the vascular structure with the vector information (e.g. Figure 6(c)). A complication of PCA, phase wrapping artefacts [14], can be clearly seen in this way (Figure 3(b)). The viewer may also visually ‘interpolate’ noisy or disjoint flow in small vessels. The results of analysis of the raw voxel data often suit graphical representation (e.g. surface of flow lines). The 3D analysis window may be used for display of such results, examples of which are discussed below.

3.3 Cut-view

The **cut-view** window is an auxiliary display below the main window control panel (Figure 1) which can be used to locally interrogate the vector/tensor data over a limited area. The cut-view can be used in both qualitative and quantitative analysis. Applications of cut-viewing

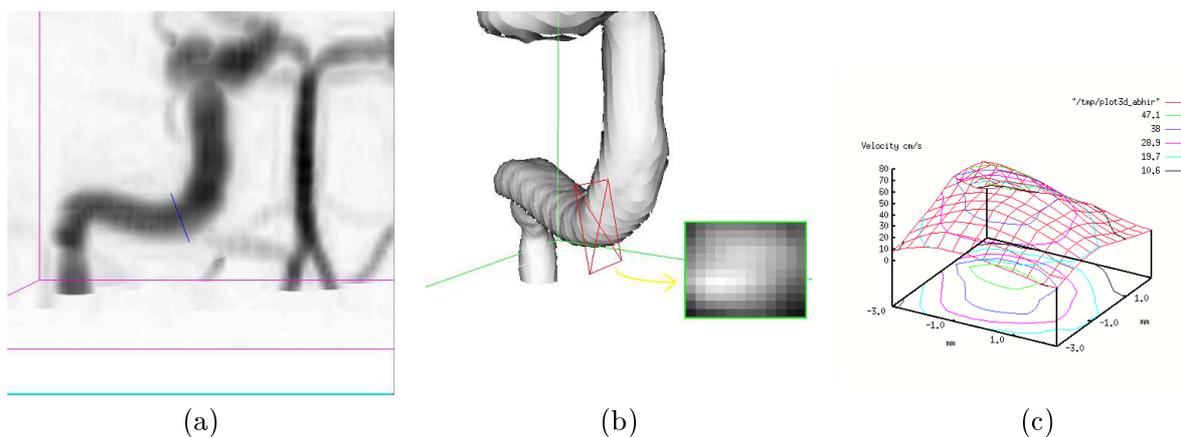


Figure 4: The user can perform a virtual ‘cut’ of a vessel. The cut plane is always orthogonal to the view-plane. Its orientation is selected by orienting a rubber-banded line as shown in (a). (b) Shows the resulting cut plane highlighted in pink and the speed data across the cut (inset view (b)). Angiotool can plot the profile of the calibrated speed across the vessel as shown in (c).

include displaying the speed profile across a vessel at a chosen point (Figure 4(c)), and precisely positioning the 3D cursor e.g. for seeding particles in flow simulations as described below. In relation to diffusion data, projections of the eigenvector directions onto or through the plane may be displayed. The cut-plane is always perpendicular to the view of the render display and centred at the most recently selected 3D cursor position. A virtual knife (a rubber banded line that follows the cursor) is used to define the orientation of the cut-plane simply by defining the ends of the cut with mouse clicks (Figure 4(a)). Alternatively, Angiotool can auto-select the cut-plane orientation to be perpendicular to the average local velocity.

4 Analysis tools

Angiotool’s analysis tools can be categorised as static, where a single, often quantitative result is output, or dynamic, where step by step interaction produces qualitative results. If the static techniques answer the question ‘what is’, the dynamic techniques, by simulation, try to answer the question ‘what if’.

4.1 Static analyses

Angiotool has limited flow quantification features. For any 3D point selected in the render window, the speed and velocity estimates derived from raw images are shown in the scrolling text panel at the bottom of the main window. This type of output could be used in assessing whether an apparent vascular stenosis seen on the MIP views is exerting a haemodynamic effect - i.e. is the blood flow velocity significantly higher through the perceived narrowing than in the segments on either side. By defining points along the vessel, a series of speed estimates and positions are transferred to the 2D plot function for display, or to processing routines for estimation of the extent of narrowing. This longitudinal information complements the cross-sectional results obtained with the cut-view function above. Applying basic calculations to the data, the cut-view itself has been extended to estimates peak and mean velocities (in cm/s), flow (in ml/s) and proportional shear stress in a vessel cross-section. Flow is estimated only over samples where the speed is greater than 20% of the peak in the defined cross-section.

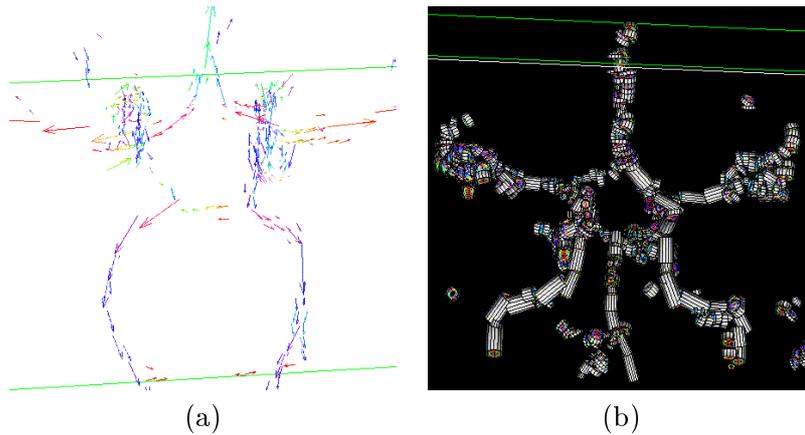


Figure 5: Summary displays of velocity information produced by a multiresolution averaging process [15]. (a) Centre line and direction information of main local features. (b) Estimates of vessel diameters are displayed as barrel motifs.

Many image processing techniques whose endpoints are data reduction and extraction have outputs suitable for graphical representation. Segmentation for example is a common step in establishing automating the detection of vessel narrowing and in defining the flow boundaries for computation fluid dynamics studies. Elsewhere we have reported on the use of the raw velocity information in a multiresolution averaging process both for segmentation and to produce summary centreline and bounding shell estimates for vessel segments [15]. These lend themselves to rendering with line and barrel motifs respectively as illustrated in Figure 5 (a) and (b) respectively.

4.2 Dynamic analyses

The potential for dynamic interaction with vector and tensor field data in Angiotool is exemplified by its facility for streamline tracking. The estimation of streamlines is a means of approximating the path of blood flowing along a vessel [3] or putative connectivity of neurons by following the anisotropic component of diffusion along a nerve bundle [7, 8, 10]. Streamline tracking is initiated by setting a number of particles (or seeds) into the velocity vector (diffusion tensor) field. For each seed Angiotool determines subsequent positions on the basis of the local velocity (diffusion anisotropy). The resulting track is displayed as a 3D curve in the analysis window.

The tracking process is local and relies solely on the velocity data at each point using a physical space, point tracking algorithm. Any time *dependent* flow stream can be expressed by the ordinary differential equation for the change in position \vec{r} given the local velocity \vec{v} :

$$\frac{d\vec{r}}{dt} = \vec{v}(\vec{r}(t), t) \quad (1)$$

giving,

$$\vec{r}(t + \delta t) = \vec{r} + \int_t^{t+\delta t} \vec{v}(\vec{r}(t), t) dt. \quad (2)$$

Since our PCA data is a time averaged velocity field $\vec{v}(\vec{r})$, a simple 1st order Euler integration can solve for the integral on the rhs without having to resort to an elaborate multi-stage numerical integration (e.g. [16]) i.e.

$$\vec{r}(t + \delta t) = \vec{r} + \delta t \vec{v} \quad (3)$$

The incremental step δt can be chosen to give displacements of the order of the minimum voxel dimensionality e.g. 0.5 or 1mm. Linear interpolation is used to estimate the velocity at the real co-ordinates \vec{r} given the discretely sampled data \vec{v}_i . For display purposes a spline curve is interpolated through the set of points $\vec{r}(t)$. The tracking process is terminated if the speed falls below a threshold value set by the user, or the track exits the data volume. Blood-flow tracking is illustrated in Figure 1 for a normal subject, and Figure 6(a)-(c) for a patient with a giant cerebral aneurysm.

With diffusion tensor data, the tracking is performed on a derived vector field such as the principal eigenvector which represents the local anisotropy modulated by a scalar. The diffusion tensor can be expressed as a linear sum of the outer product of its principle components or using tensors of rank $1 \leq i \leq 3$, \mathbf{T}_i which have a geometric interpretation depending on the relative 3D ‘shape’ of the local diffusion coefficient:

$$\mathbf{T} = (\lambda_1 - \lambda_2)\mathbf{T}_1 + (\lambda_2 - \lambda_3)\mathbf{T}_2 + \lambda_3\mathbf{T}_3 \quad (4)$$

where λ_i are the eigenvalues, and $\mathbf{T}_i = \sum_j^i \vec{e}_j \vec{e}_j^T$ are the sum of the outer product of eigenvectors respectively [8]. For the tensor tracking, we simply use the rank 1 (line) case where diffusion has taken place anisotropically in the direction \vec{e}_1 :

$$\vec{v}(\vec{r}) = f(\vec{r})\vec{e}_1(\vec{r}) \quad (5)$$

and f can be any appropriate scalar measure e.g. $\lambda_1 / \sum_i^3 \lambda_i$. Our tracking process does not deal optimally with points where fibre tracts meet or cross (where rank 2 and rank 3 tensors are required). A more sophisticated approach (see for example [8]) will be needed to better handle this situations.

The basic tracking process can be modified by reversing the time steps to track ‘backwards’ through the flow i.e. $\delta t \rightarrow -\delta t$. This is particularly useful to identify potential feeding vessels to arterio-venous malformations and aneurysms (see Figure 6(b)). For tracking white matter tracts, both forward and backward steps are taken from the seed point. This overcomes the arbitrary choice between eigenvectors and their negatives when decomposing the tensor field. An example of tracking within diffusion tensor data is shown in Figure 6(d).

Multiple seed points can also be tracked simultaneously (Figure 6(c) and (d)): either by seeding a 26-neighbourhood of voxels around the seed point, or all points in the data volume at a specified step interval can be seeded. The latter modification is quite slow, but the results can give a better impression of flow connectivity than the static vector display alone. For display purposes, arrow heads can be added to start/ends of the individual stream lines. For finer discrimination of flow streams, start points can be seeded with greater precision on the cut-view display window. This overcomes the possible ambiguities in depth in setting the seed position from the MIP alone.

When the user finds a view particularly useful, its parameters may also be book-marked, for easy return after subsequent manipulations. In similar fashion, the current content of the analysis window can be buffered at the press of button. These buffers enable successive analyses to be permanently recorded, compared, or animated. As the user has full, interactive control of the viewpoint (magnification and orientation), this provides a rudimentary means of generating and storing a sequence of changing viewpoints (fly-through). We have also found the buffering useful for recording successive stages in the streamline tracking process. A stack of display buffers is maintained by Angiotool which the user can cycle through (replay) using the movie control buttons (e.g. see example frames shown in Figure 6(e)).

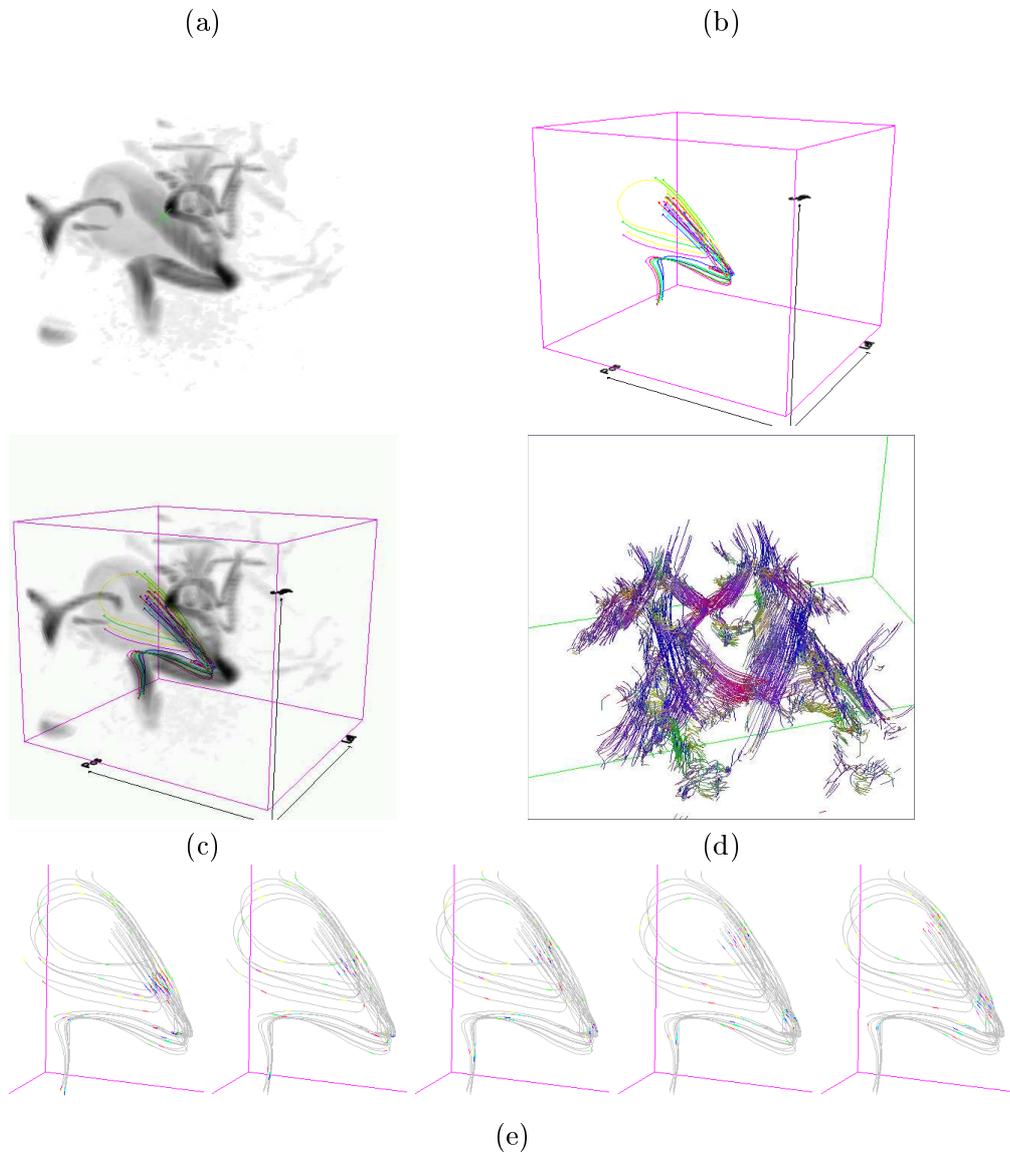


Figure 6: (a)-(c): Flow tracking experiments to determine the neck of a giant cerebral aneurysm: (a) User selects starting point(s) in MIP view. (b) By reversing the direction of flow from the seed point, blood flow is followed from within the aneurysm back through the neck and into the feeding artery, precisely locating the neck (near bifurcation off right internal carotid and middle-cerebral arteries). Further experiments with flow tracked forward in time reveals the flow vortex within the aneurysm. (c) Overlay of MIP with flow experiment. (d) Example of estimated white-matter tracks within a diffusion image from a central slice across the brain. (e) Frames from animation of aneurysm tracking experiment showing movement of virtual particles through vessels. This animation is normally rendered in 3D and user can interactively alter their viewpoint.

5 Conclusions

There has been growing interest in the development of image processing and analysis methods for angiographic and diffusion data. Examples of methods range from simply enhancing the display of curvilinear structures [4, 5, 6] to analysing flow and connectivity e.g. [2, 4, 3, 16, 10]. Inevitably, most such methods are specific to the nature of the data: time of flight or phase contrast PCA, and, on the whole, have been implemented for the purposes of algorithm development making them awkward or unsuitable for clinical use. Consideration of the end user is a recognised factor in making post-processing methods accepted and routinely used.

Angiotool provides functionality that focuses on a set of accountable and responsive operations for data exploration. The current implementation is built upon open standards technologies: GUI toolkits using X and Motif and 3D graphics using OpenGL. Although Angiotool is not as generalised and extensible a framework as other packages (e.g. AVS, Analyze, IDL, VTK) offer, we believe that its specialised nature makes it easier to use in the clinical environment and better suited to particular diagnostic tasks at hand. Its strength lies in the design being tailored to the requirements of its expert users: radiologists and surgeons. Where the experts demand up-and-coming processing, analysis and display algorithms can be appended to the existing features of Angiotool without substantive changes to the look-and-feel of the GUI.

Several aspects of the GUI and rendering capabilities could be undoubtedly improved upon. The organisation of the command syntax (layout of file menus and control panels) has not been particularly well optimised and there is no built-in way to undo, log, and record operations, capabilities which are important in a clinical setting. The ray-cast rendering engine could be augmented with volume rendering. A function which would be of use when studying groups of patients is a way to catalogue quantitative and qualitative results against the geometry of the patient's anatomy (e.g. the vasculature). We are currently in the process of incorporating structured data recording methods. Also, while the clinical interest driving Angiotool's development has largely focussed PCA and diffusion tensor imaging, other vector and tensor fields such as temperature gradients, and deformation are already being mapped with MRI. Nevertheless, we believe that the presented GUI model could usefully form the basis of other clinical applications of this type, where the need for efficiency in data visualisation and interrogation processes remains.

Acknowledgements

The authors would like to gratefully acknowledge the contributions of both clinical and computer science colleagues at the Surgical Planning Laboratory, Harvard Medical School, Boston and from the Division of Radiological Sciences, Kings College Medical School, Guy's Hospital, London. Notably: Drs Westin, Nakajima (MD), Cox (MD) and Profs Kikinis (MD), Jolesz (MD) and Hawkes. Thanks also to Dr Neil Roberts and Tom Barrick (MRI Analysis Centre, University of Liverpool) for providing example DTI data.

References

- [1] P. R. Moran, R. A. Moran, and N. Karstaedt. Verification and evaluation of internal flow and motion. True magnetic resonance imaging by the phase gradient modulation method. *Radiology*, 154(2):433–441, 1985.

- [2] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis. Three-dimensional multi-scale line filter for segmentation and visualisation of curvilinear structures in medical images. *Medical Image Analysis*, 2(2):143–168, 1998.
- [3] M. H. Buonocore. Algorithms for Improving Calculated Streamlines in 3-D Phase Contrast Angiography. *Magnetic Resonance in Medicine*, 31(1):22–30, 1994.
- [4] Thomas M. Koller. *From Data to Information: Segmentation, Description and Analysis of the Cerebral Vasculature*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 1995.
- [5] D. Wilson, J. A. Noble, D. Royston, and J. Byrne. Automatically finding optimal working projections for the endovascular coiling of intracranial aneurysms. In *Proc. of MICCAI Lecture Notes in Computer Science*, volume 1496, pages 814–821, 1998.
- [6] L. M. Lorigo, O. Faugeras, W. E. L. Grimson, R. Keriven, R. Kikinis, A. Nabavia, and C-F. Westin. Codimension-two geodesic active contours for mra segmentation. In *Proc. of Intl. Conf. on Information Processing in Medical Imaging*, 1999.
- [7] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance, Series B*, pages 209–219, 1996.
- [8] C-F. Westin, S. E. Maier, B. Khidhir, P. Everet, and F. A. Jolesz and R. Kikinis. Image Processing for Diffusion Tensor Magnetic Resonance Imaging. In *Proc. of MICCAI'99*, pages 441–452, 1999.
- [9] D. Weinstein G. Kindlemann and D. Hart. Strategies for direct volume rendering of diffusion tensor fields. *IEEE Trans on Visualization and Comp. Graph.*, 6(2):124–138, 2000.
- [10] C. Poupon, C. A. Clark, V. Frouin, D. LeBihan, I. Bloch, and J-F. Mangin. Inferring the Brain Connectivity from MR Diffusion Tensor Data. In *Proc. of MICCAI'99*, pages 453–462, 1999.
- [11] S. Peled, R. Kikinis H. Gudbjartsson, C-F. Westin, and F. A. Jolesz. Magnetic Resonance Imaging shows Orientation and Asymmetry of White Matter Tracts. *Brain Research*, 780(1):27–33, 1998.
- [12] J. D. Foley, V. L. Wallace, and P. Chan. The Human Factors of Computer Graphics Interaction Techniques. In J. Preece and L. Keller, editors, *Human-Computer Interaction*. Prentice Hall and Open University, 1990.
- [13] P. Lacroute. Analysis of a Parallel Volume Rendering System Based on the Shear-Warp Factorization. *IEEE Transactions on Visualisation and Computer Graphics*, 2(3):218–231, 1996.
- [14] A. Bhalerao, C-F. Westin, and R. Kikinis. Unwrapping phase in 3D MR phase contrast angiograms. In *Proc. of CVRMed-MRCAS'97 Lecture Notes in Computer Science*, volume 1205, pages 193–202, 1997.
- [15] P. E. Summers, A. H. Bhalerao, and D. J. Hawkes. Multiresolution, Model-Based Segmentation of MR Angiograms. *Journal of Magnetic Resonance Imaging*, 7:950–957, 1997.
- [16] D. Knight and G. Mallinson. Visualizing Unstructured Flow Data Using Dual Stream Functions. *IEEE Trans. on Visualization and Computer Graphics*, 2(4), 1996.