Visualization of cerebral blood vessels

Anna Puig Puig

1 Introduction

The computer assisted visualization of anatomical structures reconstructed from a set of 2D images is becoming an increasingly relevant feature of the medical diagnosis and the surgical simulations. In particular, the reconstruction and visualization of blood vessels provides better means of diagnosing and treating vascular pathologies. In this report, the visualization of cerebral blood vessels is analyzed and a new visualization method is proposed which enhances the blood vessels while keeping a lower computational cost than other existing methods.

First the specific problems of the medical diagnosis of vascular structures are presented and the requirements of the visualization are derived (section 2). The existing blood vessel visualization techniques are analyzed in relation to the requirements in order to detect their limitations (section 3). It is concluded that a symbolic model of the vascular map is needed providing both surface and volume data which therefore enables several visualization strategies (section 4).

Particularly, the shaded visualization of the volume data of the blood vessels provides high quality images. In section 5 the most suitable shading model to outline the vessel features is analyzed. Next, in order to reduce the cost of the visualization (section 6) three strategies are used: first a pre-process is performed in which the visualization parameters that are independent from the viewpoint are computed, next the number of samples to be projected is reduced and finally advantage is taken of the coherence between successive samples projection. Finally, in section 7, some simulations are performed.

2 Analysis of the problem

From a medical point of view, the vascular system is a tree of multiple branches that are blood vessels. Without loss of generality, each vessel can be considered as a cylindrical structure with a circular o elliptical cross-section ([FR92]). The cerebral vascular map is a connected structure although some malformations may produce isolated vascular structures.

Two types of vessels anomalies exist that should be detected in the analysis of the vascular data because they may produce severe, even mortal accidents. On one hand, local enlargements of the vessels sections, i.e. aneurysms and venoma, which can produce haemorrhages, and, on the other hand, local narrowing of the vessel sections which may lead to non-haemorrhagic accidents, i.e.thrombosis and embolisms. These vascular diseases are relatively frequent: they show up in more than 5% of the older than 65 population.

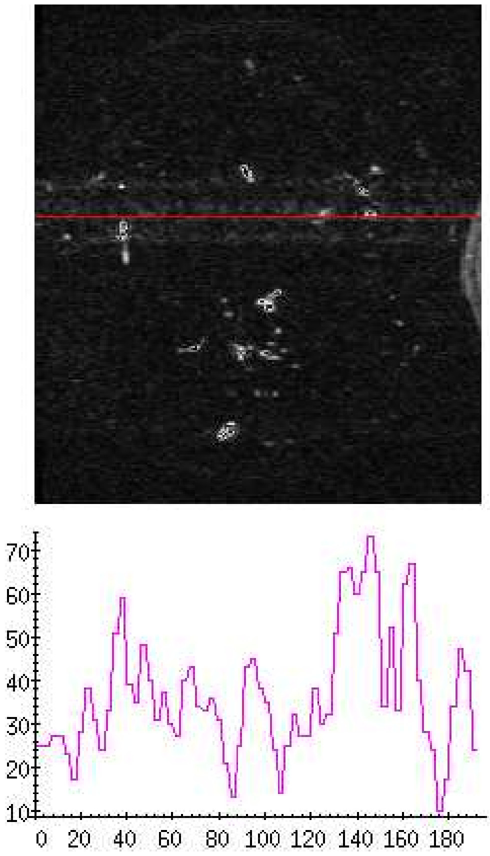

Figure 1: MRA slice and associated histogram.

The detection of zones of the vascular map having a potential risk of accident is currently based on the analysis of Magnetic Resonance Angiography (MRA) images which register at each point the density of moving hydrogen molecules ([Hor97]).

These data are generally sets of around one hundred 256x256 or 512x512 images with an approximate resolution of 1mm and 12 bits of intensity that represents the sampled density. These images have low contrast and they have noise in high intensity values because of the magnetical susceptibility to artifacts, the presence of heterogeneity in the RF fields and due to variations in the receiver antenna. In addition, if high intensity vessels exist whose section is smaller than the space between successive samples, the signal is reduced by the neighboring intensity, and the contrast is lower. Furthermore, in turbulent zones or low blood velocity regions, the signal may be lost. Therefore, the intensity values in samples which are interior to a vessel do not belong to pre-fixed range. They are simply higher than the neighboring values. However, not all the local maxima correspond to vessels, some of them may be produced by a high intensity noise, as shown in figure 1.

Because of characteristics of the MRA data above described and as the blood vessels are very narrow tubular structures, the perception of the vascular map by direct observation of the data is very difficult and it does not enable the automatic detection of anomalies. A computer assisted processing of the images and a posterior visualization are required in order to realize diagnoses ([VVB+94]).

Synthesizing, the visualizations should have the following features:

- Perception of the vascular topology, of its characteristics (regular sections, stenoses , aneurysms and branching) as well as the relationships between these features.
- Three-dimensional perception of the structures
- Perception of the vascular surface
- Perception of the type of material that the blood vessels are made of.
- Interactivity
- Zooming of regions of interest
- Integration of the vascular data and the stationary tissue.

3 Previous work

The visualization process is composed of different stages from the data acquisition to the final image: identification of the vascular structure, reconstruction of the topological and geometrical structure and visualization itself, i.e. shading and projection. These two latter computations depend on the internal representation scheme which is used.

The identification of the blood vessels may be done in a process previous to the visualization (segmentation and transfer function edition) or during the visualization itself, while classifying voxels, i.e. assigning to them color and opacity values. Several authors have developed segmentation algorithms suitable for blood vessels. These methods are able to detect narrow tubular structures with low contrast and small objects causing signal attenuation. They manage low noise signals and heterogeneity in the intensity taking into account the photometric properties of the data ([HTL+89], [CLKJ90]), the shape of the blood vessels ([GKJ90], [EDKS94]) and the continuity of the vascular map ([GKJ90]), [VVM+93]).

The reconstruction of the vascular structure consists of obtaining a symbolic representation model for the blood vessels that contains semantical information, such as the location of aneurysms and stenoses, as well as topological and geometrical information. Generally, the symbolic model is a logical hierarchical structure which, according to the vessels features, organizes superficial ([BGSC85], [EDKS94]) or volumetric information ([ZJP94]).

The identification of the vessels as well as the reconstruction of the symbolic model are optional steps of the visualization process.

Independently from the existence of the two former stages, the visualization methods can be classified into two groups depending on the representation model that they use:

- Methods which extract the surface of the vessels and visualize it *(surface rendering)*
- Methods which visualize directly the data *(volume rendering)*

3.1 Surface Rendering

The *surface rendering* methods construct a surface model from the data, either manually ([BGSC85]) or automatically ([EDKS94]). The use of these methods with blood vessels data requires a high quality binary segmentation preprocess, given the complexity of the vascular structure. Once the model has been reconstructed, any conventional surface rendering method can be used in the visualization, generally, hardware driven Z-buffer with Gouraud or Phong shading.

The reconstruction of a surface can be done directly from the data or from a symbolic model. In the former case, the Marching Cubes technique[LC87] is the most widely used. It approximates the boundary surface of the vessels with triangle patches. Because of the topology of the vascular map, the resulting model is made of a large number of tiny faces. This model is valid for large diameter vessels because a cross section of the surface occupies different voxels and it is possible to interpolate it. In small vessels, on the contrary, the continuity of the surface may not be guaranteed, and if their section is smaller than one voxel it is even impossible to reconstruct it (see figure 2). In addition, it should be taken into account that the surfaces in the data are from the blood flow: the surface extracted is the inner vessel wall which is equivalent to the blood surface. Therefore it is smooth, at least geometrically differentiable. Moreover, it should be flexible in order to enable local deformations such as aneurysms and stenoses. This is the reason why the trilinear interpolation used in Marching Cubes is not always acceptable for blood vessels surface extraction.

Another method used in the visualization of surfaces is the Dividing Cubes ([LC91]), which can be considered as a variation of the former algorithm that obtains a model of points and normals associated to each point. This model, because of the topology of the vessels is composed of a very large number of points and presents the same inadequacy problems than the Marching Cubes approach.

In general, if the segmentation process is good enough and the reconstruction of the surface has a low error in relation to the real surface of the vessels, then the surface rendering methods give high quality realistic images and enable a good perception of the vascular topology. The surface models are however very complex and memory expensive. In addition, because of the interpolation that they do, they sometimes do not fit to the true physical object. Some symbolic models simplify the surface model taking into account some physical features of

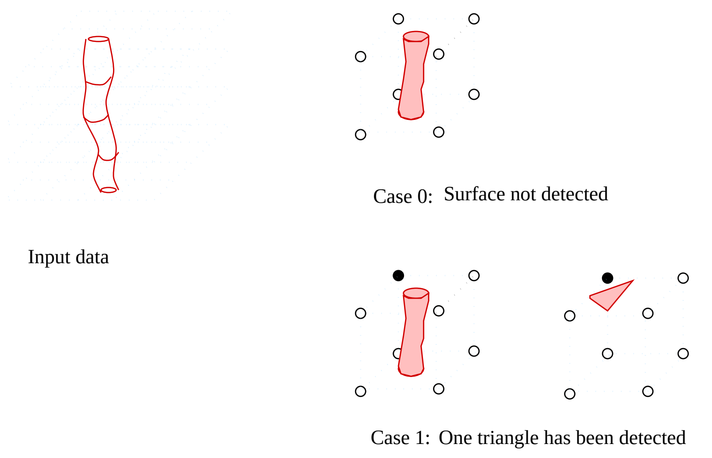

Figure 2: Problems of the Marching Cubes Technique for Blood Vessels Surface Extraction.

the real objects such as the fact that blood vessels have approximately circular or elliptical sections. In [EDKS94] a symbolic model is constructed from which a surface model made of biquadratic splines with rectangular basis is extracted. This model is suitable for vessels with diameters larger than 1mm. In addition the symbolic models allow the explicit definition of the interrelationships between the different features.

According to the requirements defined in section 2, a 3D conventional visualization of the surface model such a Z-Buffer with Phong shading, allow interactive rates, zooming of regions of interests and the use of textures to show the material type. However special algorithms should be developed in order to integrate the stationary tissue in the visualization.

3.2 Volume Rendering

The other technique used in the visualization of the data is *volume rendering*. In this strategy, the set of data is directly projected and some classification functions are used in order to assign to each voxel an opacity and a color that allow a visual identification of the surface and the material type. This avoids the problems of the binary segmentation of the surface extraction strategy. In order to compute the shading, different behavior models of the object with light are considered, including emission, absorption and single and multiple scattering.

The volume rendering methods visualize directly the voxels model, either slice-to-slice [CP92], or with splatting [Wes90], or by ray-casting [Lev88]. Although several methods based on a scan-line strategy have been published [SMA96], they are most of all used for tetrahedral lattices.

The kernel of all the projection strategies is the computation of the radiance of each pixel of the final image. This radiance can be evaluated by integrating the transport equation along the rays of vision associated to each pixel. This requires the computation of several shading parameters at different samples along the rays, generally more than one per voxel. The main difference between the Splatting and the RayCasting approaches is that the former one integrates simultaneous all the rays while the latter one integrates each ray separately.

The computation of the shading parameters at sample locations, either voxels or points of the rays, is based on the fact that the property values at these samples define the type of structure that occupies the voxel and they can be computed using transfer functions. Thus the different ranges of property values should correspond to different structures. However, in MRA data, this assumption is not true, because MRA values of voxels interior to a blood vessel do not belong to a specific range, they are simply higher than the values of the neighboring voxels. Thus the integration along a ray in MRA data would compose both stationary and vascular intensity, giving a lower contrast to the vessels than the original slices.

This is the reason why, in MRA data, rather than an integration, it is generally preferred to simply project the maximum intensity value of the data in the corresponding pixel. This method is known as Maximum Intensity Projection (MIP). The MIP method is simple, it gives at first glance good image quality but it does not allow real-time visualizations. Although a maximum intensity projection could be obtained either in a projective or a ray-casting approach, it is the former one which is generally used, and, in the bibliography, the expression MIP refers to a RayCasting with MIP projection.

Several accelerations of the algorithm have been proposed based on the minimizing the number of interpolation of property values along the rays GS95]. The main drawback of the MIP are inherent to the method itself, because any voxel interior to the blood vessel that is not enough contrasted with the stationary tissue around it will not be reflected in the visualization; and any voxel with local noise can be visualized as belonging to vessels. For this, there may be losses of small, low-contrasted vessels; the diameter of the vessels may be artificially reduced because in MRA data the vessels edges are less contrasted , and if the images are very noisy, it is the maximum noise which is visualized. In addition, this technique does not provide three-dimensional perception. Therefore, the topology of the vessels may be misunderstood (see figure 3), there are no depth clues and no reference of the stationary tissue. Because of all the problems the MIP, different alternate visualization techniques have been proposed.

In order to provide depth clues to MIP projections, sequences of images from different viewpoints are often computed and small animations are projected. However a static image is sometimes preferable in order to analyze anatomical details. A different approach is to use stereoscopic displays. Finally, it is also possible to shade the point of maximum intensity according to its depth [HBP+90].

In order to reference the blood vessels in relation to the stationary tissue [HTL+89] proposes to perform a mixed ray-casting that detects the maximum value along the ray and then integrates the intensity from this point, along the ray of vision.

In Table 1 the methods analyzed are classified according to the requirements exposed above. Synthesizing, on one hand the MIP technique does not keep the topology of the vessels, nor the tridimensionality of the data, and it can produce artifacts. On the other hand, it allows reasonably fast visualization integrated

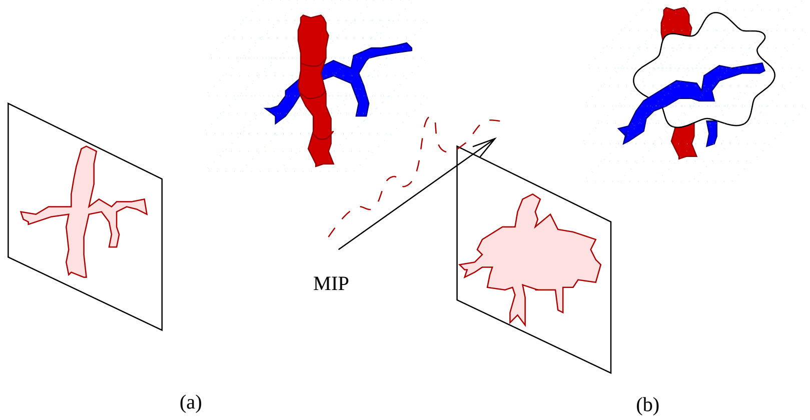

Figure 3: Problems of the MIP technique: (a) in projected vessels the topological relationship between them is not detected (they can be either a real branching or the crossing of the projection of two independent vessels) (b) the maximum noise is projected (the white cloud is noise of maximum value).

	Surfaces	Volume	
	Marching/Dividing	MIP	Composition
Vascular Topology	NO	NO	YES
3D	YES	NO	YES
Vascular surface	Segmentation	YES	classification
Material type	texture	shading	shading
Interaction	YES	fast	NO
Zoomings	YES	YES	YES
Tissue integration	hybrid	combinat	YES

Table 1: Analysis of the methods of blood vessels visualization according to the physicians requirements.

with the stationary tissue and with zooming of the areas of interest. The volume rendering techniques are in general slower than the MIP ones and than the surface rendering approach, but they give 3D images than do not change the vascular topology and that emphasize the perception of the vessels on the basis of the classification results. With symbolic models, as the reconstruction of the surface is enabled while keeping the voxel data, the volume rendering approach allows to obtain better quality integrated images. Thus their optimization is a challenging alternative to the other methods.

In [Pui98] a symbolic model of the vascular map allowing the representation of its topology, its surface and its volume information is presented. Next, this model is briefly described and several strategies for its visualization are analyzed.

4 Visualizations of a symbolic model of the cerebral vascular map

4.1 Description of the model

The symbolic model of the vascular structure is a graph such that its nodes store the branching of the vessels and its edges keep up the information of vessels segments between branching. Each edge of the graph is subdivided into as many segments as vascular features exist in the vessel. These features have been identified as aneurysms, stenoses and normal vessel sections. Each element of this graph stores both surface and volume data.

The superficial information is represented as a union of Generalized Cylinders. For each vessel, the set of sections across the central axis of the vascular structure, i. e. the skeleton curve, is stored. Each section represents the contour curve at different levels of detail, from a simple circle or ellipse to a spline. This information is stored in the edges of the graph model.

The volume data of the segmented blood vessel voxel model are run-length encoded at two levels. The segmentation stage maintains the values of the internal properties of the vascular structure while losses only the external voxels values. A new property is added to the input voxel model such that it labels each voxel with the identifier of the graph element to which the voxel belongs. In the first level, run-length encodes directly this new property on a compressed voxel model. In the second level, the run-length divides the input voxel model on the basis on the symbolic graph elements and for each element of the graph it encodes the voxel submodel according to the input values of the real vascular properties. As this volume model stores pointers to the surface model and as it is partially on the graph, the relationships between volume and surface data are guaranteed. In addition, the compressed voxel model keeps up the spatial information which allows to traverse the model back-to-front or front-to-back as a non-compressed voxel model.

4.2 Visualizations

The advantage of the symbolic model is that it enables several possible visualization strategies exploiting thus the benefits of the different approaches surveyed above:

- Direct visualization of the symbolic model in order to allow a quick glance of the vascular topology and to detect possible anomalies,
- Visualization of the vascular boundary which can be hardware driven and thus provides interactivity,
- Three types of volume rendering:
  - Slice-to-slice visualization which allow a fast perception of the volumetric properties,
  - MIP projection, which gives images familiar to the physicians, such as X-Rays images, or

  - Shaded visualization which provides high quality images by keeping the topology of the vessels and the tridimensionality of the data,

- Visualization of the hybrid model which renders both surface and volume data and
- Visualization of the vascular model integrated with the stationary tissue.

The two former visualizations can be performed with conventional surface visualization methods depending on the interpretation of the inter-contour associated information. To render this surface, conventional techniques can be used.

The slice-to-slice projection and the MIP projection are realized with the methods presented in section 3.

The shaded volume rendering as well as the hybrid and integrated visualizations must be analyzed specifically in order to guarantee the physicians requirements. First, the optical model most suitable to emphasize the blood vessels surface while keeping a low error is studied. Next, an optimization of the visualization of the vascular map is proposed, which allows the integration with the vascular tissue and enabling zooming. The efficiency of the visualization is obtained by isolating in a pre-process the shading computations that are independent from the viewer and by exploiting the coherence between successive voxels in the projection.

5 Optical model of the blood vessels

5.1 Light integration

The light transport through heterogeneous materials has been exhaustively studied ([Bli82], [Max95]). Any participating media has a set of optical properties that enable the emission of light, the absorption and the scattering.

The general *Time-Invariant Gray Radiance Equation* [Gla95] expresses the radiance at a point r and a direction w as the sum of radiance of the first surface point s hit by a ray traced from r in direction w, attenuated by the participating media between r and s plus the radiance at any volumetric point $r - xw$ of the ray, attenuated between $r - xw$ and r:

$$L(r, w) \quad = \quad \beta(r, s)\Gamma(s, \omega) + \int_r^s \beta(r - x\omega, x) Q(r - x\omega, \omega)\delta x \qquad (1)$$

where $L(r, w)$ is the radiance at s in direction w, $Q(r - xw, w)$ is the radiance at a distance of x from r on the ray and $\beta(r', r)$ is the attenuation function between any point r' and r.

The attenuation function also called *path absorption* modelizes amount of light absorbed and out-scattered between r' and r. It can be expressed as:

$$\beta(r', r) \quad = \quad e^{\int_r^{r'} \Theta_\lambda(t)\delta t} \qquad (2)$$

being $\beta(r, r')$ the optical distance between r' and r.

The application of this general model to medical data may be done according to several hypothesis: either all the anatomical structures are considered to behave as fuzzy participating media, with no boundary, or one or more particular structure in the volume are considered as solid discretized objects whose boundary should be rendered as a surface model. In this latter case, the anatomical structures around and inside the surface of interest are treated as participating media.

For vascular models, according to the requirements, in the visualization, the surrounding and internal volume of the vessels should be shaded as participating media whereas the vascular wall should be shaded as a surface. Therefore, the rays crossing the vascular model should be integrated through the volume until the surface is hit (see Figure 4). If there is not intersection of a ray with the surface then a background radiance value will simply substitute the surface radiance.

Rendering the vascular data involves therefore solving Equation 1. Several simplifications, usual for medical data can be done. First, the scattering onto volume particles is not considered as it leads to very complex images, difficult to understand in a diagnosis context. On the contrary, the model most used supposes only emission and scattering. Second, rather than computing radiances, light intensities I_λ, for each wave length λ, are simply treated.

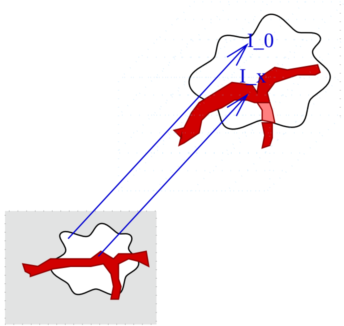

Figure 4: Modeling of the light along the ray.

Although the analytic form is the most correct way to evaluate the integral, the numerical approximation with the Riemann sums is enough:

$$I_\lambda(D) = I_{\lambda_x} \prod_{i=1}^{n} t_{\lambda_i} + \sum_{i=1}^{n} g_{\lambda_i} \prod_{j=i+1}^{n} t_{\lambda_j} \tag{3}$$

where D is the observer's position, $t_{\lambda_i} = e^{-\Theta(i\Delta x)\Delta x}$ and $g_{\lambda_i} = g(i\Delta x)$.

It can be considered that Δx is the length of the segments in which the integral is discretized (i.e. D/n where n is the number of segments). Several methods exist that sample the ray at regular intervals ([Lev88]) and adaptively to a surface in the volume ([Lev90b]). In both strategies, however some data may be lost if voxels hit by the ray are missed. This produces object-space aliasing in the final image.

As a consequence of the discrete partition of space in the voxel model, this sums can be computed by sampling the intervals according to the voxels structure (see figure 5). Each of segment of the discretization of the sum 3 can be computed as the ray path through each voxel. The intervals interior to a voxel may be characterized according to different properties: probability of hitting a vascular surface crossing the voxel, type of variation of the property in the voxel (constant, linear, ...) and according to the behavior of light inside each voxel (existence and type of light attenuation, ...).

5.2 Shading inside the voxels

Thus, two types of voxels are considered: those which are crossed by the surface, called *boundary voxels*, and those that completely interior or exterior to the vessels *fuzzy voxels*. The illumination model is adaptive to the information associated to each type of voxel.

Fuzzy voxels

Inside the fuzzy voxels, the property values, and therefore the optical behavior may be considered as constant or it may vary linearly. Besides, the attenuation of light may be omitted or on the contrary, it may be computed according to the length of the ray inside the voxel.

Depending on the level of realism desired, different methods have been developed in order to approximate the intensity and the opacity accumulation along a ray inside a fuzzy voxel ([WG91]) from averaging the rear and front intensities to integrating the exponential form of the attenuation. In the former strategy, if the opacity is computed as $min(1, \Omega * d)$, then the opacity and intensity of a voxel are:

$$g_{\lambda_i} = \frac{g_{\lambda_i}}{\Theta} * t_{\lambda_i} \tag{4}$$

$$t_{\lambda_i} = min(1, \Omega * d) \tag{5}$$

This method is called C*D in [WG91], because the intensity is simply the depth by the if $t_{\lambda_i} < 1$. This approximation gives good results if the property value and the attenuation inside a voxel are more or less constant. The visualization of large volume data sets with this method is hardly distinguishable from other more complex shading techniques. Therefore, it is usually acceptable in most cerebral vascular voxel models [1].

In order to compute the intensity and the opacity of each voxel, i.e. g_{λ_i} and t_{λ_i}, the scalar values of the properties of the original data associated to each

[1] The basic assumption of the C*D model is that the attenuation is from voxel to voxel but not inside a voxel, thus visualizing a voxel with an edge length 1000 will not produce the same result than a projecting 1000x1000x1000 homogeneous voxels.

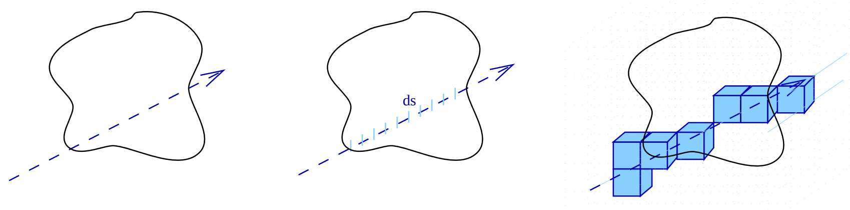

Figure 5: Different successive approximations of the integration of light along a ray.

point are used as entries to the transfer functions from which the R, G, B, α values of the sample are retrieved. The transfer functions used are based on segmentation models of the MRA data proposed in [VVB+94]. They are stored as look-up tables.[2]

Boundary voxels

In vascular boundary voxels, being the surface opaque, the integration of the rays inside the voxel stops when the surface is hit. Thus, the lighting model in boundary voxels composes the intensity along the ray inside the voxel up to the surface with the shading at the surface (see figure 6). It should be noted however that not all the rays crossing a boundary voxel hit the surface (case (a) and case (b)), so some rays crossing a boundary voxel may be integrated further through other real voxels inside the volume.

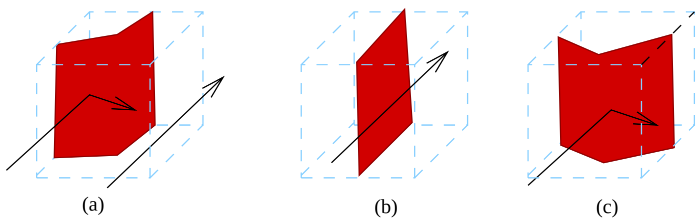

Figure 6: Contribution of the boundary voxels.

There are several approaches to compute the contribution of the volume inside the voxel in front of the vessel surface [Lev90b], although for blood vessels, given the occupancy ratio between volume and surface and according to the input data, this segment of integration can be simply omitted.

In general the optical behavior of a surface is composed by scattering, transmission and emission. For vascular data, the surface is supposed to be opaque, so that the transmission is not modelized. In addition, the vessels do not behave as luminaires so there is no emission. Moreover, the surface scattering is purely diffuse and perfectly specular, as glossy reflections do not seem to be proper for vascular surfaces. Finally, the surface scattering model is local, as vessels are not supposed to reflect radiance one to each other. It is supposed that an

[2] The transfer functions are customized to the specific type of analyzed data. For blood vessels, they adapt to MRA data.

external light source iluminates the surface from the observer's position. It is as if the viewer has a lamp to explore the data. The atenuation of this external source through the volume can be modelized empirically or applying in the reverse sense the atenuation along the volume to the surface.

Therefore the local illumination of the surface can be computed by simply applying an extension of Phong shading ([BT75]):

$$I_{\lambda_x} = I_{\lambda_a} k_{\lambda_a} + f_{att} I_{\lambda_p} [k_{\lambda_d} (\bar{N}\bar{L}) + k_{\lambda_s} (\bar{N}\bar{H})^n] \tag{6}$$

where k_{λ_a}, k_{λ_d}, k_{λ_s} are the coefficients of ambient, diffuse and specular reflection, respectively. n is the exponent of specular reflection $\bar{N}$ is the surface normal, f_{att} is the attenuation factor, L is the light direction, that, as mentioned above, is substituted by the viewer direction and I_{λ_p} is the external light intensity. Finally H is the direction of maximum highlight.

The computation of the surface normal cannot be done directly from the data value. Generally, with medical data, it is approximated by the gradient of the property value in relation to a given neighborhood. In most applications, where the surfaces are simple enough, this approximation suffices. However it is unacceptable for the vascular surface which is a narrow tubular structure of large curvature. In particular, in [HBP$^+$90] it has been shown that the gradient may give image quality for large vessels but that it produces numerous artifacts in small vessels. In order to realize an exact computation, it is necessary to start from the hybrid model that contains the volumetric information as well as surface data. This model enables to access efficiently to the surface that passes through a voxel and to compute its location and orientation [Pui98].

6 Optimization of the shaded visualization

In the visualization process, to analyze each sample, two types of computations must be performed: the projection of the geometry and the shading of the sample. The efficiency of the visualization depends on the cost of these operations. Next, an optimization of the visualization based on the C*D shading is proposed, that guarantees the quality of the final image, while keeping a low computational cost. It is applicable to the two classical volume rendering approaches (projective and RayCasting). This optimization may be applied to the visualization of any other structure that could be shaded according to the C*D model, although it is especially suitable for blood vessels.

6.1 The Weight Matrix

With the model C*D (see the equation 5), in order to compute the contribution of a voxel to any image pixel, the shading value associated to the voxel according to its property value must be weighted with the depth of ray of vision corresponding to the pixel.

Assuming parallel projections, the projected area of a voxel can be applied as a footprint to any other voxel by only translating it. An unique reference vertex for each voxel should be projected in order to determine the location of the voxel projection into the image buffer.

Therefore, the Weight Matrix, $W(v)$, of a voxel v is defined as:

$$W(v) = \{w(i,j),\ 0 \leq i < with(v),\ 0 \leq j < height(v)\ / \\ W(i,j) = Nlength(\sigma_{ij})\} \quad (7)$$

where σ_{ij} is the intersection segment between the voxel v and the sight line, λ_{ij}, passing through the pixel (xmin+i, ymin+j) and $Nlength(\sigma_{ij})$ is the normalized length of the intersection segments.

This matrix keeps up the projection geometry and the distance contribution at any pixel of any voxel.

On parallel projections and on the basis of the spatial coherence between voxels, the Weight Matrix of a voxel can be used as a footprint of any other voxel applying only a translation. The coherence between voxels relies onto two basic voxel properties. First, any voxel of the input model is homeomorphous to any other one, next, the length of the intersection segment of a 3D line, λ_{ij}, and a voxel v is equal to the length of the intersection between a line homeomorphous to λ_{ij} and a voxel homeomorphous to v.

However, the use of the Weight Matrix directly as a footprint presents some problems due to the discretization of the image space: the translation is discrete, shared faces exists between neighbor voxels which must be projected only once and the projected edges can be shortened in the rasterization. To solve these problems, the definition of the Weight Matrix is extended to the concept of the Generalized Weight Matrix (GW).

The discrete translation problem stems from the fact that discrete projections of two homeomorphous voxels are identical while the length values of any pair of intersections of two corresponding rays on the two voxels may be different. In other words, the value of the weight of the two projections can vary depending on the spatial position of the voxels (see Figure 7).

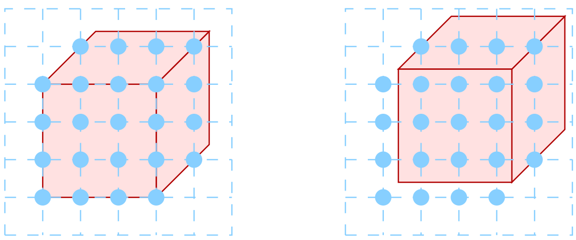

Figure 7: Discrete translation problem.

To solve this problem, an averaged Weight Matrix is computed that stores at each pixel the average depth that each pixel may have.

On the other hand, the faces shared by any pair of neighbor voxels must not be projected twice. To avoid this, the averaged Weight Matrix is computed as the rasterization of the semiopened silhouette of the projected pattern voxel (see Figure 8).

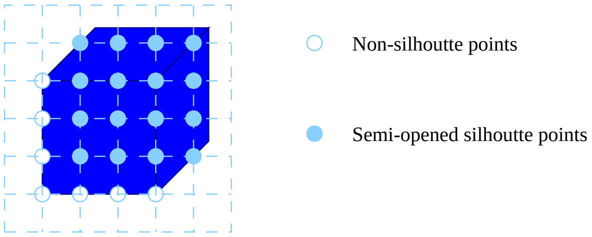

Figure 8: Semiopened silhouettes.

Summarizing, the Generalized Weight Matrix, $GW(v)$, of a voxel v is defined as:

$$GW(v) = w(i,j), \; 0 \leq i < with(v), \; 0 \leq j < height(v)$$
$$w(i,j) = \begin{cases} Alength(\sigma_{ij}) & if \; (i,j) \in Silhouette(v) \\ 0 & if \; (i,j) \notin Silhouette(v) \end{cases} \tag{8}$$

where $Alength(\sigma_{ij})$ is the average length of the intersecting segments of the line sight with the all the virtual cubes that have an identical raster projection than the voxel v. The intersection of the ray with all the voxels that have the same discrete projection than the voxel v. $Silhouette(v)$ is the semiopened silhouette of the projected voxel v.

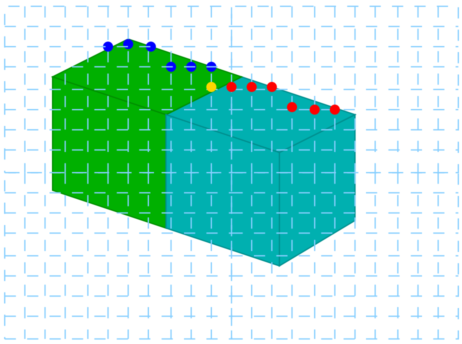

Figure 9: Problem of location of the projection.

Finally, from any viewpoint position, the width and the height of any voxel projection are not integer values in most cases, thus, when they are discretized, these measures are shortened. Because of this, after projecting several successive voxels, the projection of the next voxel may be shifted of one pixel in any image direction. This shifting can occur repetidly along the voxels projection. The moment in which this error may occur depends on the accumulated error of the rasterization of the previous voxels (see Figure 9).

The combination of the possible positions into the two image directions gives a total of sixteen possible weight matrix, according to the rounding error ([TPN95]).

Thus, 16 footprints are computed in a pre-process and they are used in the projection algorithm according to the location of the voxel in the raster image.

6.2 Projective methods

To apply the Generalized Weight Matrix in the projective methods, it is only necessary to project the reference point of any voxel in order to define the location of the voxel projection in the image buffer. Moreover, the location of a new voxel can be calculated from the projection of the previous voxel using only integer arithmetic. This avoids to transform the reference at each voxel. The computation of the shading at any voxel is trivial from the stored information stored in the associated Generalized Weight Matrix.

The ordered Back-To-Front (BTF) projection of the voxel model for a given viewing direction selects as the primary traversal plane, the coordinate voxel plane with maximum projected area. This allows to compute the largest weight matrices and thus to maximize the benefits of the use of the Generalized Weight Matrix. In addition, the direction of integration has minimum depth inside the voxels and thus the error of the C*D composition is minimized.

6.3 RayCasting methods

The RayCasting methods are based on the composition of the sampled intensities along each parallel ray which starts at the pixels of the image buffer and traverses the volume data according to a viewing direction.

The Weight Matrix optimizes the RayCasting methods in the computation of the intersection of the ray with the voxels in different ways.

First, the calculation of the intersection between the ray and the voxels can be realized incrementally starting from an initial discretization of the ray according to the voxel mesh and from the codification of each subarea of the voxels through which the ray is going. The discretizations of the next rays can be computed to add and to remove voxels. If, in addition, the discretization stores the shade value computed at each voxel, the per-voxel shading is performed only once per voxel as in the projective approach, by opposite to conventional RayCasting methods that compute it as many times as rays hit a voxel. This is a different approach than those which obtain the shade value in a pre-process and store it as a property value, as these latter ones multiply at least by 4 the initial input voxel size ([Lev90a]).

On the other hand, [YK92] proposes the template RayCasting technique in order to compute in incremental mean the discrete ray that involves a 2D final rotation of the calculated image with the consequent smooth error. The basic idea of the proposed optimization is to take advantage of the Weight Matrix, above defined, in order to reduce the time of the voxel sampling along a ray. [MY96] proposes a similar optimization but the voxel projection used is based on the splatting technique ([Wes90]).

According to the defined projection, in order to identify which of the 16 possible footprints should be applied it is only need to know the reference vertex of the voxel and the ray equation.

Each footprint is divided into subareas. Each subarea is a set of pixels that share the same pair of intersected faces of the voxel (see Figure 10). These subareas have been defined in [WG91]. Therefore, with the reference vertex and a ray equation, the relative entry on the footprint and the subarea can be determined.

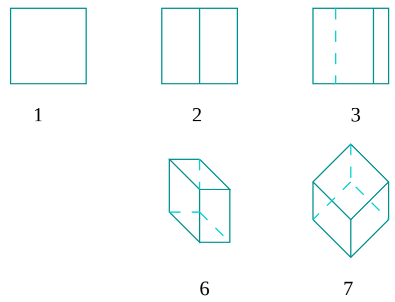

Figure 10: Weight Matrix subarea.

For voxel projections of more than one pixel, the set of rays that share the same discrete ray is defined as **tunnel**. Then, the light contribution of the whole set can be computed in a one-pass strategy with the knowledge of the voxels of the discrete ray and their associated footprints. Likewise, the coherence between two consecutive tunnels can be exploited to calculate incrementally the next tunnel and the associated set of footprints (see Figure 11).

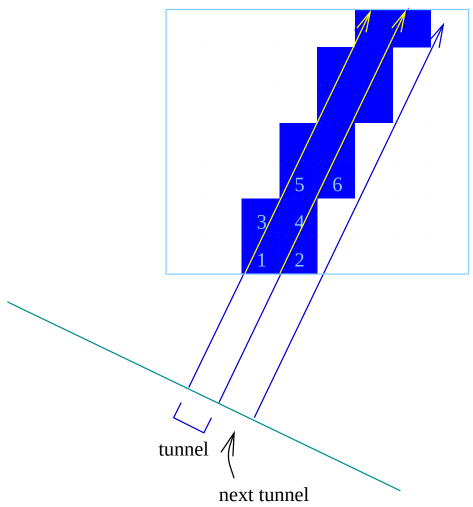

Figure 11: Tunnel definition and coherency between tunnels.

6.4 Cost of the optimization

The proposed optimization decreases the cost of the projection of each voxel depending on the size, in pixels, of its projection.

Let n^3 to be the number of the voxels of the input data and let $px \times py$ to be the number of pixels onto which the bounding box of the voxel set is projected given a particular viewpoint, then the ratio voxel-pixel is defined as:

$$rat_x = \frac{px}{px_{max} - px_{min}} \tag{9}$$

$$rat_y = \frac{py}{py_{max} - py_{min}} \tag{10}$$

where,

$$px_{max} = \forall i : 0 \leq i < 8 : max(o_{x_i}) \tag{11}$$

$$py_{max} = \forall i : 0 \leq i < 8 : max(o_{y_i}) \tag{12}$$

$$px_{min} = \forall i : 0 \leq i < 8 : min(o_{x_i}) \tag{13}$$

$$py_{min} = \forall i : 0 \leq i < 8 : min(o_{y_i}) \tag{14}$$

and $(o_{x_i}, o_{y_i}, o_{z_i})$ are the limit points of the bounding box of the voxels set and $(v_{x_i}, v_{y_i}, v_{z_i})$ are their corresponding observer,

$$\forall i : 0 \leq i < 8 : (v_{x_i}, v_{y_i}, v_{z_i}, 1).M_{TrVis} = (o_{x_i}, o_{y_i}, o_{z_i}, 1)$$

The proposed algorithm is valid for projections where the ratio pixels/voxel is larger than one. If the ratio is minor, more than one voxel may project into one pixel, so either any one of these voxels shade values is chosen to compute the pixel intensity, or an average of all them. In both cases, significant data may be lost. This is not acceptable for vascular data. Summarizing, the hypothesis of a ratio larger than one is not a restrictive limitation, on the contrary, it is a requirement for a good image quality.

In Table 2, the following costs are defined: cost of a *brute-force* projective method C_{bf}, cost of a shade projective algorithm $C * D$ proposed by [WG91] C_w and cost of the proposed optimization C_{WM}. The set of the diagrams 12 and 13 analyzes the bounds where each algorithm is optimum. In Figure 12, a set of diagrams compares the costs of these algorithms depending on the number of projected voxels with different ratios pixels-voxel (1, 5, 25, 100 and 1024).

In order to explore exhaustively the domain of the ratio, in Figure 13 the functions $C_{bf} = C_W$ i $C_{bf} = C_{WM}$ are analyzed. From this, it can be seen that the algorithm proposed by [WG91] has a lower cost than the *brute-force* method given any number of voxels and independently from the ratio value. On the contrary, the proposed optimization in relation to the *brute-force* algorithm, becomes more effective starting from a lower ratio bound larger which depends on the number of projected voxels. However, this bound is acceptable according to the parameter values of the different visualizations. Summarizing, the analysis shows that the relation between C_W and C_{WM}, it can be concluded that the optimization method has a lower cost on visualizations with a high number of voxels and ratios greater than one.

Methods	$O(n^3, rat_x, rat_y) = n^3 \times (c_{proj} + c_{shade} + c_{acc})$
Conventional C_{bf}	$n^3 \times (6 \times (c_{projpunt} + c_{win-view})+$ $(rat_x\ rat_y) \times [c_{raigvoxel} + c_{dist} + c_{c*d}]+$ $(rat_x\ rat_y) \times c_{acc}$
[WG91] C_w	$6 \times c_{projpunt} + c_{cerca}+$ $n^3 \times (14 \times c_{acces} + 8 \times c_{inter}+$ $5 \times c_{transf}) + 10 \times c_{promig} + 14 \times c_{c*d}+$ $(rat_x\ rat_y) \times [c_{inter} + c_{acc}])$
Weigth Matrix C_{WM}	$c_{WM} + c_{projpunt} + c_{win-view} + n^3 \times [(c_{trasldisc})+$ $(rat_x\ rat_y) \times [c_{accesWM} + c_{c*d}]+$ $(rat_x\ rat_y) \times c_{acc}]$

Table 2: Costs of the projective methods.

where,

$$c_{WM} = (rat_x\ rat_y) \times [c_{raigvoxel} + c_{dist} + c_{casaXY}] + 16 \times c_{casaXY}$$

The computations of these costs are based on the projection cost of every voxel, c_{proj}, on the cost of the shade calculation, c_{shade} and on the accumulation stage cost, c_{acc}.

Finally, in Table 3 the comparison of a conventional ray-casting method and the proposed extension is shown. The comparative diagrams 14 illustrate the number of the input voxels and the number of the ratio pixel-voxel value from which the optimization is suitable. In conclusion, starting from only one voxel, the optimization is a better approach than a conventional one.

In addition, in Table 4 the decrease of the costs due to the optimizations in relation to the conventional methods are written. In the diagram 15, it can be observed that the increase of the cost on projective methods is nearly constant with the number of input voxels while with RayCasting methods, the decrease of the cost grows with the number of input voxels. Therefore, the proposed optimization is most suitable on RayCasting methods than on projective ones.

Methods	$O(n^3, p_x, p_y) = p_x \times p_y \times c_{raig}$
Conventional	$p_x \times p_y \times (3\ n\ \sqrt{3}(c_{raigvoxel}+$ $+c_{proj} + c_{win-view} + c_{dist} + c_{c*d} + c_{acc})$
WeigthMatrix	$c_{discr-raig} + c_{WM}+$ $p_x \times p_y \times (3\ n\ \sqrt{3}(c_{raig-act} + c_{accesWM} + c_{c*d} + c_{acc})$

Table 3: Costs of the RayCasting methods.

Methods	$\Delta O(n^3, p_x, p_y)$
RayCasting	$p_x \times p_y \times (3\ n\ \sqrt{3}(c_{raigvoxel}+$ $+c_{proj} + c_{win-view} + c_{dist} + c_{c*d} + c_{acc})-$ $c_{discr-raig} + c_{WM}+$ $p_x \times p_y \times (3\ n\ \sqrt{3}(c_{raig-act} + c_{accesWM} + c_{c*d} + c_{acc})$
Projectives	$n^3 \times (6 \times (c_{projpunt} + c_{win-view})+$ $(rat_x\ rat_y) \times [c_{raigvoxel} + c_{dist} + c_{c*d}]+$ $(rat_x\ rat_y) \times c_{acc}-$ $c_{WM} + c_{projpunt} + c_{win-view} + n^3 \times [(c_{trasldisc})+$ $(rat_x\ rat_y) \times [c_{accesWM} + c_{c*d}]+$ $(rat_x\ rat_y) \times c_{acc}]$

Table 4: Comparison of the optimizations on projective and RayCasting methods.

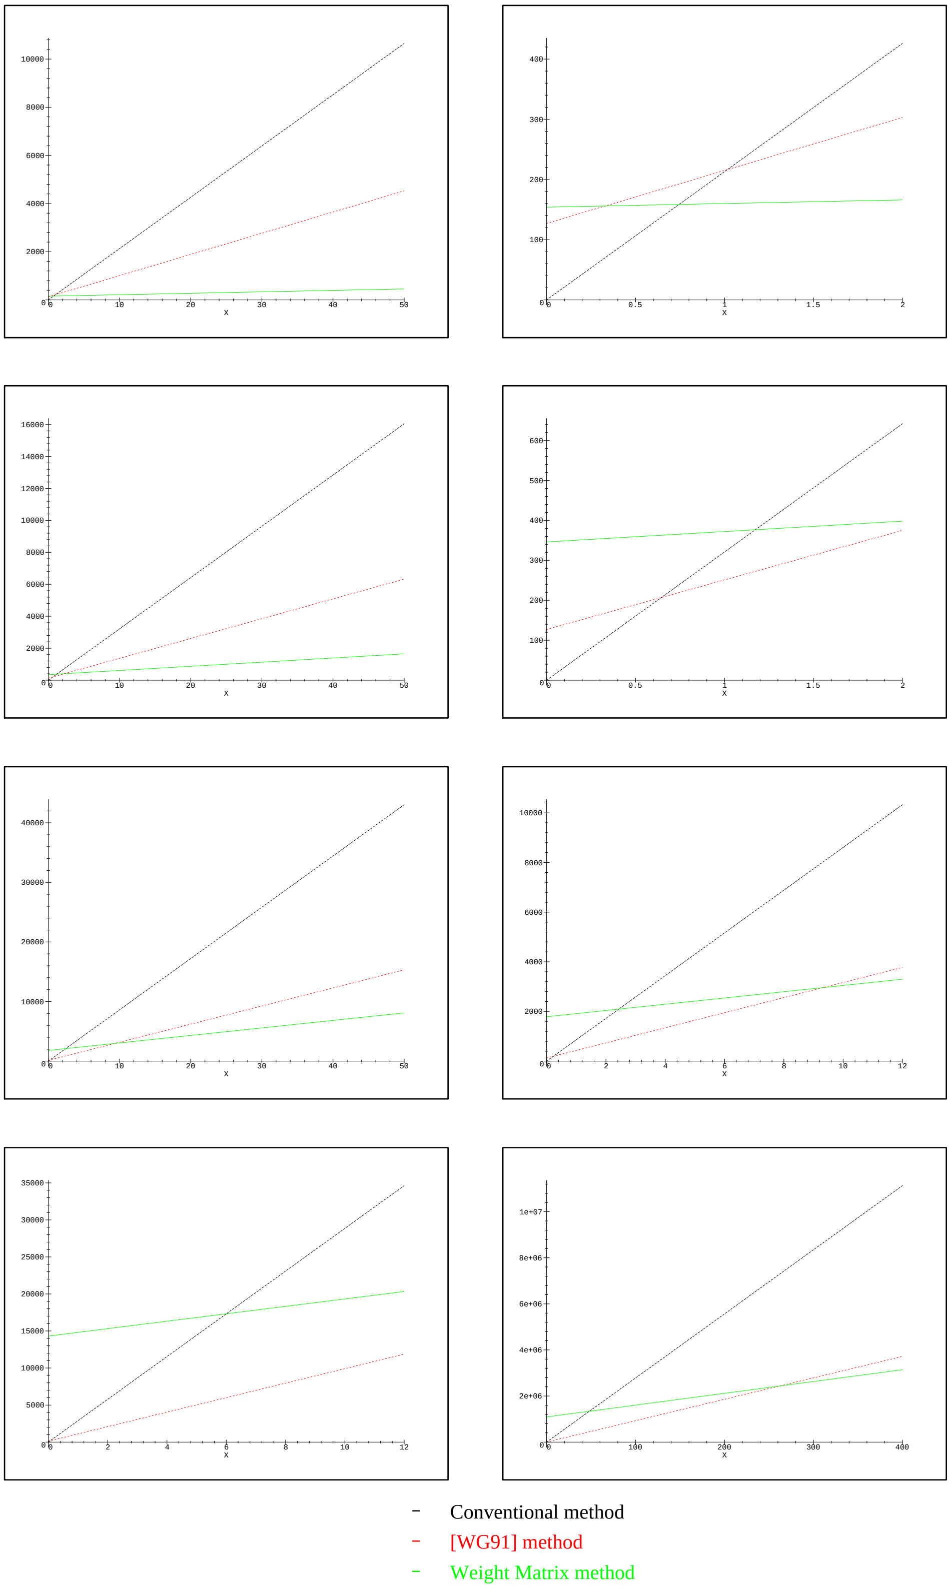

Figure 12: Comparative analysis of the costs of the projective methods with different values of ratios pixel-voxel: 1, 5, 25, 100 and 1024. The second, fourth and sixth figures are zoomings of 1, 5 and 25 samples respectively. The horizontal axis represents the number of voxels to be visualized.

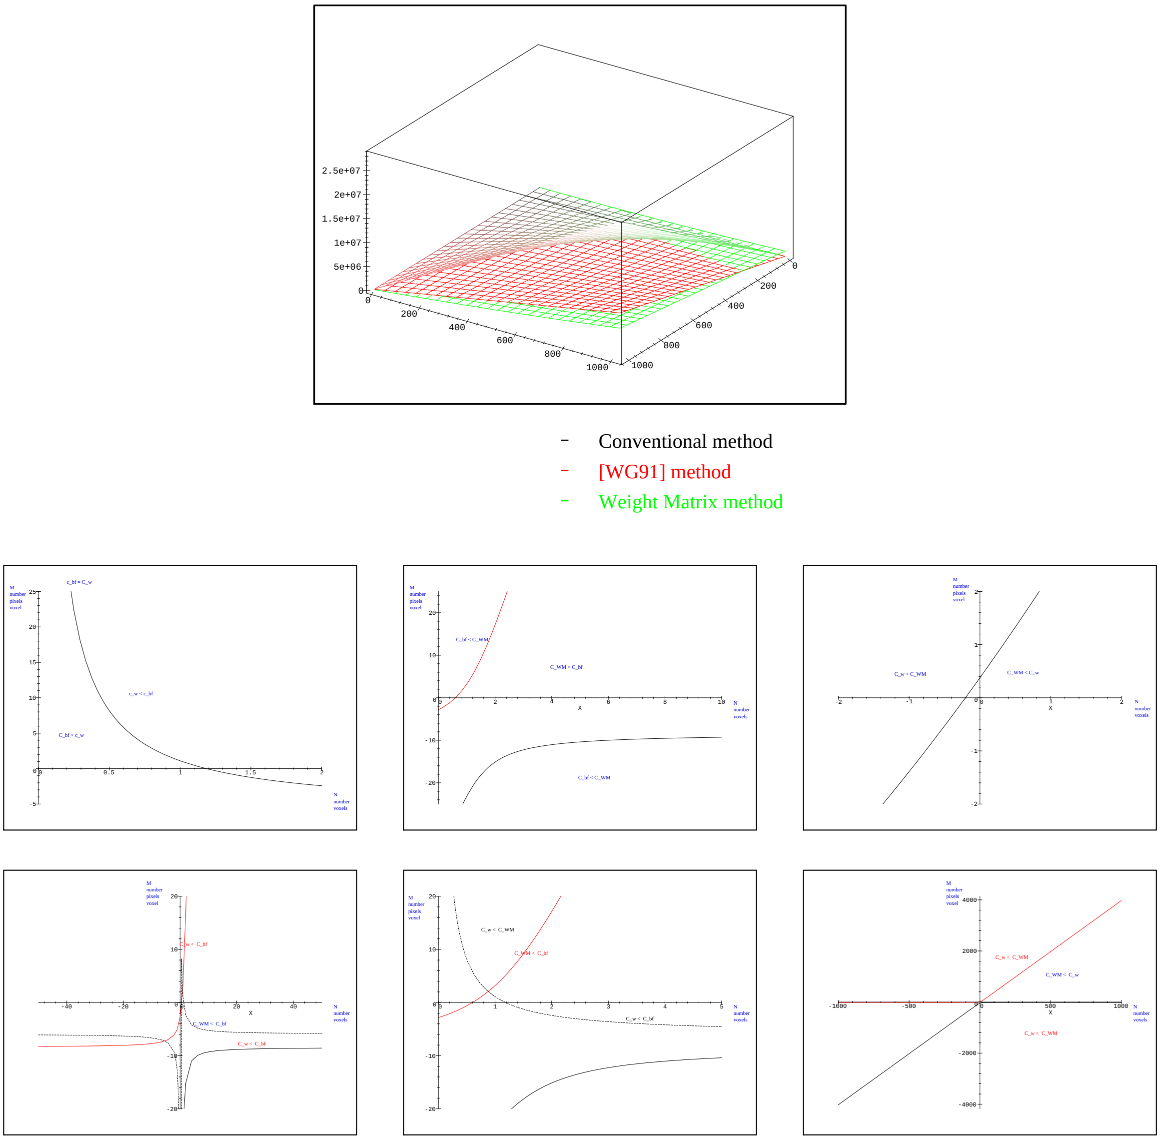

Figure 13: Comparative analysis of the projective methods.

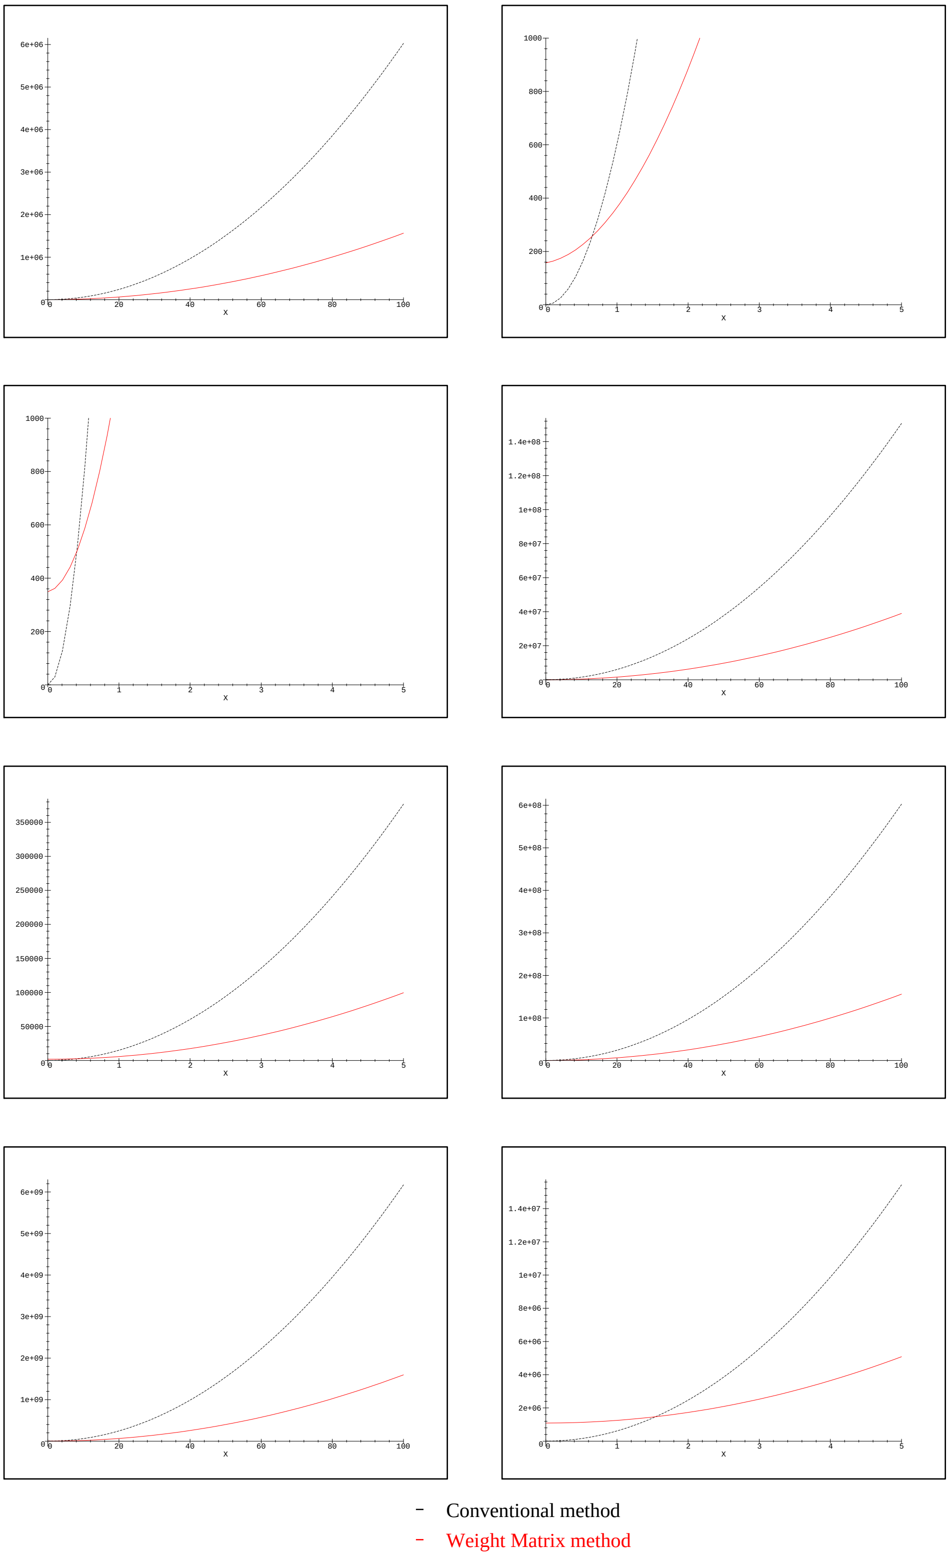

Figure 14: Comparative analysis of the RayCasting methods with different values of the ratio pixel-voxel: 1, 5, 25, 100 and 1024. The second, fifth and eighth figures are zoomings of 1, 25 and 1024 samples respectively. The horizontal axis represents the number of voxels to be visualized.

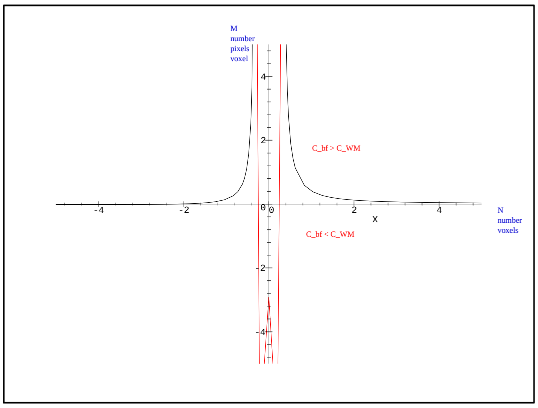

Figure 15: Comparative analysis of the RayCasting methods in relation to the number of voxels (N).

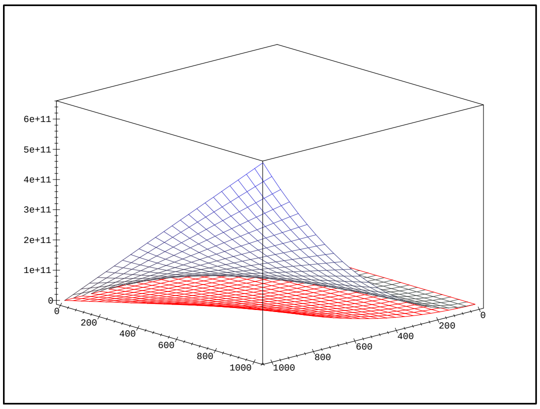

Figure 16: Comparative analysis of RayCasting methods.

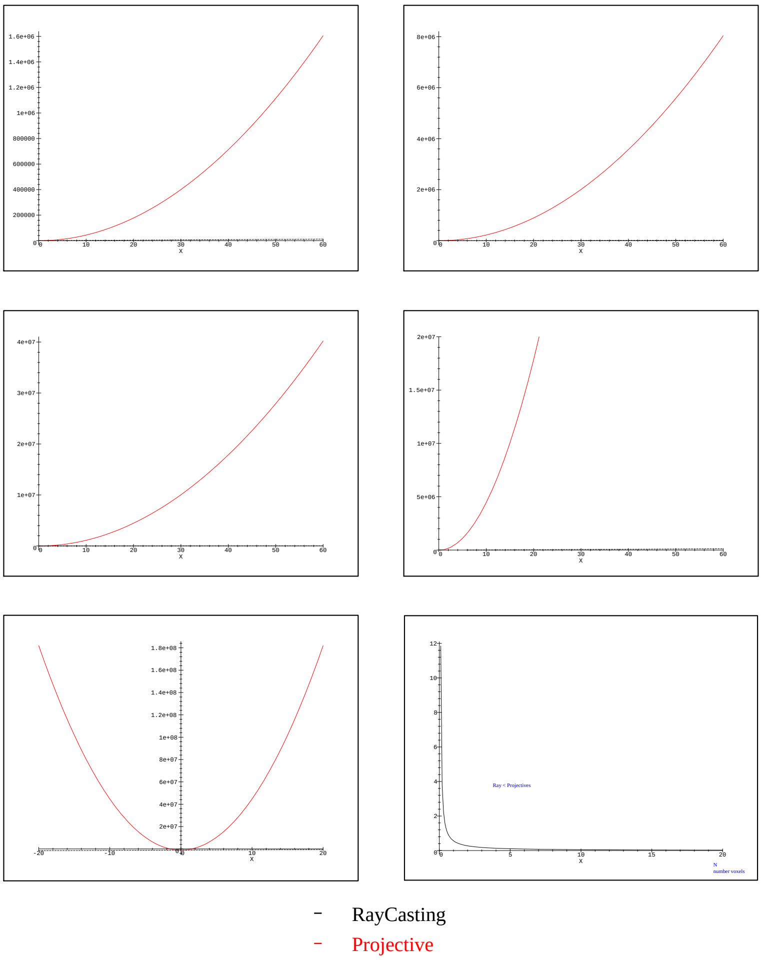

Figure 17: Comparative diagrams of the decrease of the cost on projective and RayCasting strategies with different values of the ratio pixel-voxel: 1, 5, 25, 100, 1024. The last figure is the comparative analysis between the optimizations of projective and Ray-Casting methods according to the number of voxels of the model.

7 Simulations

The symbolic model is rendered as the surface model in a simplified way (see Color Plate 18). This simplification assumes that the nodes of the symbolic graph are spheres and the edges are polylines such that their significant points have associated a contour curve approximated by a circle. The significant points of the polyline are those where the polyline has a C^1 discontinuity and those where a vascular anomaly occurs. This visualization is used to manage the user interface of the vascular model and to render the topological relations between different blood vessels.

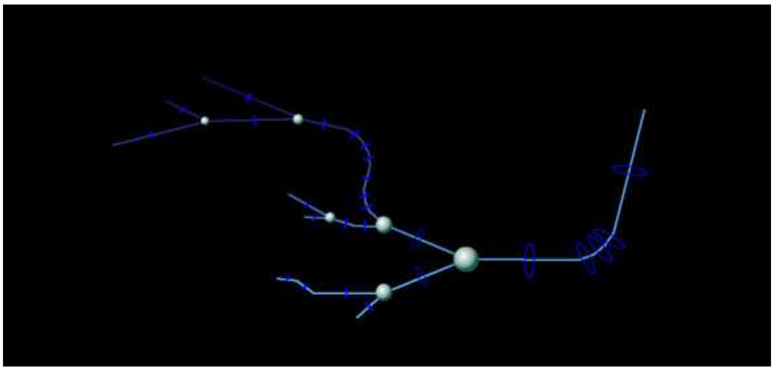

Figure 18: Symbolic Model Visualization: Polylines.

Moreover the symbolic model, or the surface model at low level of detail, can be visualized as a piping system ([JASTV93]) if the edges of the graph are interpreted as a cylinders and the nodes as spheres (see Color Plate 19). Polygonal approximations with a hardware Z-Buffer are used to visualize it. It is therefore fast and it allows navigation inside the pipes as well as interrogations on the semantics and the geometry of the model.

The visualization of the higher level of detail of the vascular surface can be performed by interpreting the inter-contour associated information as planar faces or as biquadratic splines parameterized with rectangles [EDKS94]. To render this surface (see Color Plate 20), conventional techniques can be used [FDFH93].

The visualization of volume model may be performed by slice-to-slice composition, by MIP projection and by semi-transparent shaded rendering.

The slice-to-slice composition, inspired on the BOB system [CP92], simply overlaps slices of the voxel model parallel to the coordinate plane of maximum projected area, and maps the source pixels onto target pixels with a monochrome level proportional to the property value (see Color Plate 21). It allows the composition of the images with transparencies and supports different resolutions in the voxel projection by collapsing blocks of voxels in a slice. A coherent strategy for the replication of source pixels is implemented [LTPN96], which improves the performance of the visualization. Although its quality is insufficient to meet the diagnosis requirements, this projection is fast enough to be the default visualization of the system on which the user interacts in order to select the input

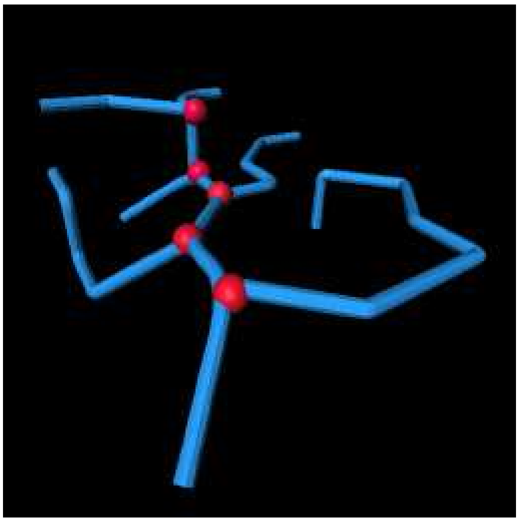

Figure 19: Symbolic Model Visualization: Pipes and spheres.

parameters of the two other voxel visualizations: clipping planes defining the region of interest of the voxel model, active ranges of the property values (i.e. values taken into account in the visualization)) and camera definition. In our experience, clinicians feel comfortable interacting on slices as in conventional 2D packages.

The MIP projection (see Color Plate 22) and the shaded rendering are executed as separate processes within the application in order to reduce waiting time.

The shaded proposed herein has been tested with medical data of the head. It is currently being tested with vascular data (see Color Plate 23). Hybrid (surface and volume) visualizations of this model and integrated visualizations with the brain are currently being implemented.

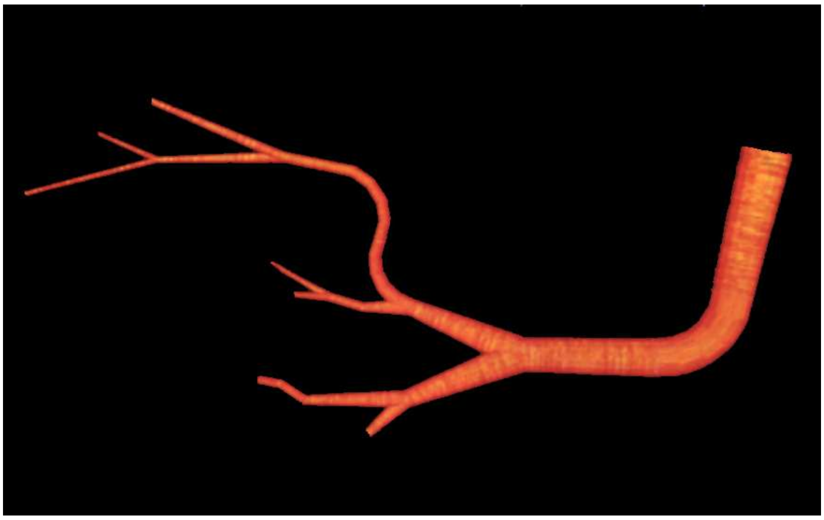

Figure 20: Surface Visualization.

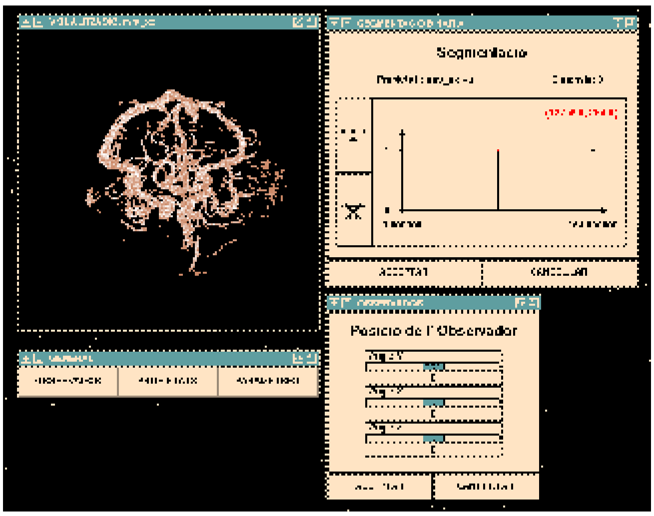

Figure 21: Slice-to-slice Volume Visualization.

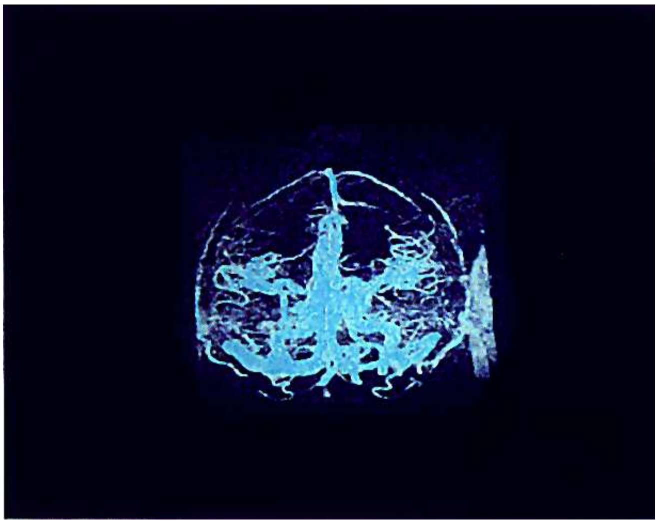

Figure 22: MIP projection Volume Visualization

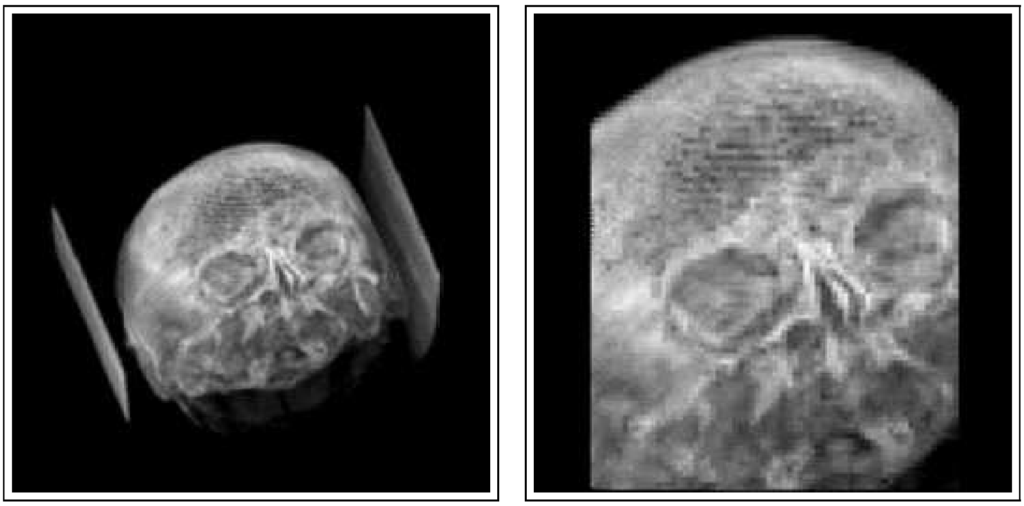

Figure 23: Skull semi-transparent shaded.

8 Evaluation

To evaluate the method proposed herein, the requirements defined in section 2 are tested:

- The perception of the vascular topology is obtained with the vascular model which isolates the features and that defines the geometry and the topology of the vessels while guaranteeing the continuity of the structure. These properties allow to enhance the abnormal features such as aneurysms or stenoses with conventional methods such as color differentiation, highlighting, squared, ...

- The tridimensional perception of the structures is achieved by rendering the volume, where they are integrated with the brain, and their surface. The proposed method visualizes directly the data and it projects the voxel information. The computed shading on the projection takes into account the depth of the intersection of the viewing ray with the voxel and it is accumulated along this ray. In order to avoid introducing noise on the data values, a voxel projection method based only on the geometry projection is used. Splatting algorithms such as [Wes90] and [LH91] which attempts to reduce the aliasing attached to the voxel representation may distort the original data.

- The perception of the vascular boundary and of the material type of the vessels is obtained by the computation of the shade from the surface of the auxiliary model.

- The interactive visualizations where the viewpoint can be positioned anywhere can be attained without any transformation of the voxel model ([Han90]) and without any transformations of the final image ([YK92]). In this way, errors on the shade computation from the successive interpolations of the voxel values are avoided.

- Zooming of the regions of interest in some visualizations are enabled for larger enough numbers of projected pixels per voxel (ratios pixels/voxel).

- The interactivity of the vascular map visualization is obtained by computing in a pre-process the Weight Matrixes associated to the particular viewpoint and by using them in the computation of the projection of each voxel.

- The integration of both stationary tissue and vascular structures is obtained to compose the computed intensities on both models.

9 Conclusions and future work

In this report, a new method of volume visualization which enhances the blood vessels and decreases the rendering time has been presented. According to the physicians requirements for the diagnosis of vascular diseases and taking into account the captured data, a suitable enhancing shading model has been defined. In addition, an optimization of the computation of the projection of the blood vessels, which is versatile and fits both the projective and the RayCasting approaches, has been proposed and evaluated.

The Extended Weight Matrix method allows to implement coherence optimizations with an almost negligible pre-process computational overhead. The algorithm is flexible enough to be adapted to a wide range of shading models.

Although more simulations should be done, the first implementation results have proved that proposed method produces interactive visualizations of the vascular model as well as of the integration with the brain. It is especially suitable when the projection of a voxel is more than one pixel. Moreover, it is suitable to the navigation inside the vascular map.

Work in progress is the application of the Extended Weight Matrix method to allow internal volume navigation on the blood vessel model, using a precalculated set of Extended Weight Matrix of a discrete set of viewpoint directions.

Other research lines on the application of the Extended Weight Matrix are:

- Adapting it to the cell-to-cell representation, i.e., to non homogeneous properties inside the voxels.
- Taking benefit from the order in which the voxels are stored in order to allow performant memory managements. This latter improvement would not guarantee a minimum C*D error, but the relation speed-up versus image quality should be analyzed.
- Implementing hardware driven projection and accumulation.

References

[BGSC85] C. Barillot, B. Gibaud, J. Scarabin, and J. Coatrieux. 3d reconstruction of cerebral blood vessels. *IEEE Computer Graphics and Applications*, 5(12):13–19, December 1985.

[Bli82] J.F. Blinn. Ligth reflection functions for simulation of clouds and dusty surfaces. *Computer Graphics*, 16(3):21–29, July 1982.

[BT75] P. Bui-Tuong. Illumination for computer generated pictures. *Communications of ACM*, 18(6):311–317, June 1975.

[CLKJ90] H.C. Cline, W.E. Lorensen, R. Kikinis, and F. Jolesz. Three-dimensional segmentation of mr images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037–1045, November/December 1990.

[CP92] K. Chin-Purcell. Bob: Brick of bytes. *AHPCRC, Minnesota SuperComputer Center (Software for Silicon Graphics)*, 1992.

[EDKS94] H.H. Ehricke, K. Donner, W. Koller, and W. Straber. Visualization of vasculature from volume data. *Computer and Graphics*, 18(3):395–406, 1994.

[FDFH93] J. Foley, A. Van Damm, S. Feiner, and J. Hughes. *Computer Graphics, Principles and Practice*. Addison Wesley, 1993.

[FR92] S. Flesia and M. Roche. 3d reconstruction of generalized cylinders from biplane projection. *Computer and Graphics*, 16:167–173, 1992.

[GKJ90] G. Gerig, R. Kikinis, and F.A. Jolesz. Image processing of routine spin-echo mr images to enhance vascular structures: Comparison with mr angiography. *3D Imaging in Medicine: Algorithms, Systems and Applications*, pages 121–132, 1990.

[Gla95] A. Glassner. *Principles of Digital Image Synthesis*. Morgan-Kaufman Series, 1995.

[Han90] P. Hanrahan. Three-pass affine transformations for volume rendering. *ACM Computer Graphics*, 24(5):71–78, November 1990.

[HBP+90] K.H. Hohne, M. Bomans, A. Pommert, M. Riemer, U. Tiede, and G. Wiebecke. Rendering tomographic volume data using: adequacy of methods for diferents modelitis and organs. *3D Imaging in Medicine*, pages 197–216, 1990.

[Hor97] J.P. Hornak. *The Basics of MRI.* http://www.cis.rit.edu/htbooks/mri/mri-main.html, 1997.

[HTL+89] X. Hu, K.K. Tan, D.N. Levin, S.G. Galhorra, C.A. Pelizzari, G.T.Y. Chen, R.N. Beck, C.T. Chen, and M.D. Cooper. Volumetric rendering of multimodality, multivariable medical imaging data. *Proceedings of Chapel Hill Workshop on Volume Visualization*, pages 45–49, 1989.

[JASTV93] R. Juan-Arinyo, LL. Solano, E. Torres, and J. Vega. Sistema dmi de modelado geométrico de tuberías. *Proceedings del 3ero Congreso Español de Informática Gráfica*, pages 223–236, 1993.

[LC87] W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM Computer Graphics*, 21(4):163–169, July 1987.

[LC91] W.E. Lorensen and H.E. Cline. Volume modeling. *Course Annotes of SIGGRAPH '91*, 7:31–43, 1991.

[Lev88] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8:29–37, May 1988.

[Lev90a] M. Levoy. Efficient ray tracing of volume data. *ACM Transactions on Graphics*, 9(3):245–261, July 1990.

[Lev90b] M. Levoy. A hybrid ray tracer for rendering polygon and volume data. *IEEE Computer Graphics and Applications*, 10(8):33–40, March 1990.

[LH91] D. Laur and P. Hanrahan. Hierarchical splatting: A progressive refinemet algorithm for volume rendering. *Computer Graphics*, 25(4):285–318, July 1991.

[LTPN96] J. López, D. Tost, A. Puig, and I. Navazo. Voldmi: An open system for volume modeling and visualization. *Computer and Graphics*, 20(5):703–712, 1996.

[Max95] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, June 1995.

[MY96] K. Mueller and R. Yagel. Fast perspective volume rendering with splatting by utilizing a ray-driven approach. *Proccedings of Visualization'96*, pages 65–72, 1996.

[Pui98] A. Puig. Cerebral blood vessels modelling. *Report LSI-98-21-R, Polytechnical University of Catalunya, Spain*, 1998.

[SMA96] C.T. Silva, J. Mitchell, and A.Kaufman. Fast rendering of irregular grids. *Proceedings of Symposium on Volume Visualization'96*, pages 15–22, 1996.

[TPN95] D. Tost, A. Puig, and I. Navazo. A volume visualization algorithm using a coherent extended weight matrix. *Computer and Graphics*, 19(1):37–45, 1995.

[VVB+94] D. Vandermeulen, R. Verbeeck, L. Berben, D. Delaere, P. Suetens, and G. Marchal. Continuous voxel classification by stochastic relaxation: Theory and application to mr imaging and mr angiography. *Internal Research Report, Medical Imaging Research Lab, ESAT, Belgium*, 1994.

[VVM+93] R. Verbeeck, D. Vandermeulen, J. Michiels, P. Suetens, G. Marchal, J. Gybels, and B. Nuttin. Computer assited stereotactic neurosurgery. *Image and Vision Computing Volume*, pages 468–485, October 1993.

[Wes90] L. Westover. Footprint evaluation for volume rendering. *Computer Graphics*, 24, August 1990.

[WG91] J. Wilhems and A. Van Gelder. A coherent projection approach for direct volume rendering. *ACM Computer Graphics*, 25(4), July 1991.

[YK92] R. Yagel and A. Kaufman. Template-based volume viewing. *Proceedings of Eurographics 1992*, 11(3):153–167, 1992.

[ZJP94] C. Zahlten, H. Jurgens, and H.O. Peitgen. Reconstruction of branching blood vessels from ct-data. *Proceedins of Rostock*, 1994.