Journal of Universal Computer Science, vol. 21, no. 7 (2015), 935-958
submitted: 4/8/14, accepted: 26/6/15, appeared: 1/7/15

Non-Photorealistic Rendering of Neural Cells from their Morphological Description

Angela Mendoza
(Rey Juan Carlos University, Madrid, Spain
angela.mendoza@urjc.es)

Susana Mata
(Rey Juan Carlos University, Madrid, Spain
susana.mata@urjc.es)

Luis Pastor
(Rey Juan Carlos University, Madrid, Spain
luis.pastor@urjc.es)

Abstract: Gaining a better understanding of the human brain continues to be one of the greatest and most elusive of challenges. Its extreme complexity can only be addressed through the coordinated and collaborative work of researchers from a range of disciplines. 3D visualization has proven to be a useful tool for simplifying the analysis of complex systems, where gaining meaningful understanding from unstructured raw data is almost impossible, such as in the case of the brain. This paper presents a novel approach for visualizing neurons directly from the morphological descriptions extracted by neuroscience laboratories, pursuing two goals: improving the readability of complex neuronal scenarios and avoiding the need to store 3D models of the intricate geometry of neurons, since such models are demanding of computer resources.
The proposed rendering method involves illustration techniques that facilitate the visual analysis of dense neural scenes. The work presented here brings the field of neuroscience and the benefits of 3D visualization environments closer together, increasing the interpretability of massive neural scenarios through visual inspection. A preliminary user study has proven the utility of the proposed rendering techniques for the visual exploration of dense neuronal scenes. The feasibility of parallelizing the implemented algorithms has also been assessed, representing a further step towards interactive illustrative visualization of neuronal forests.

Keywords: 3D Visualization techniques, neuronal data visualization, illustrative rendering.
Categories: I.3.0, I.3.3, I.3.4, I.3.8

1 Introduction

The need to analyse complex systems is present in a wide variety of disciplines ranging from science to security, from economics to neuroscience. In the case of neuroscience, there are two reasons behind the huge increase in data complexity expected over the next decade: first, the availability of powerful microscopes using technologies such as the FIB-SEM CrossBeam workstations [Casares and Gnauck, 2009], [CarlZeiss, 2014], which are allowing researchers to drastically reduce the time required to gather experimental data (from months to hours). The second reason is the development of ambitious research programs such as the Blue Brain [Markram, 2006], Obama's BRAIN initiative

[Wolf, 2013] and the Human Brain Project HBP [Mcardle, 2014], which are encouraging research teams from different disciplines to join forces to advance towards the understanding of the brain.

Current neuroscience comprises many disciplines with different specific goals. The study of neuron morphology deals with the analysis of the shape and structure of cells and the impact of these features on the cells functionality. 3D visualization has much to offer in this regard, since researchers need to detect aspects related to variations in neuron morphology across different specimens, tests or brain regions, or neural patterns in large scenarios. The development of new rendering strategies for these complex neuronal scenes can facilitate the analysis and interpretation of images, increasing the quality and productivity of the researchers.

Over the last few decades, different rendering methods have been developed, focused mainly on either increasing realism or reducing computation times. Consequently, most of the contributions were devoted to the generation of high quality, photorealistic renderings [Debevec, 2008], [Lehtinen et al. 2012], and [Laine et al. 2010] or to the design of simplified techniques that generate images of good visual quality at interactive rates [Beyer et al. 2014], [Nichols and Wyman, 2010].

However, neuronal data present some peculiarities that warrant special attention. First of all, the nature and amount of data question whether representing the neurons with polygonal meshes is appropriate for visualization and for storage purposes. Secondly, even when depicting only a small fraction of the real number of cells, the complexity of a neuronal scenario is such that its visual analysis is difficult in the extreme. This paper deals about the visualization of neuronal data, focusing on two important aspects that simplify the visual analysis process: rendering data directly from the original descriptions generated by neuroscientists and designing rendering techniques, which simplify the visual analysis process. More specifically, the main contributions of this paper are:

- Design of a new rendering technique taking the morphological descriptions directly generated by the neuroscience community as input.
- Incorporation of expressive illustrative techniques into the rendering process in order to increase the visual interpretability of dense scenes.
- Parallelization of the rendering algorithms in order to assess the feasibility of running the proposed method on high-performance architectures.

The rest of the paper is organized as follows: section 2 is a brief introduction about basic concepts of neurons as well as a short overview of some previous work related to visualization techniques and rendering methods. Section 3 describes the proposed approach, and section 4 shows some experimental results. Finally, the conclusions and proposed future work are presented in section 5.

2 Background

2.1 Neurons

There are currently many software tools that can help with the extraction of the 3D shape of neurons from a stack of microscopy images. The automation of this process involves a segmentation step and the computation of a surface that approximates the cell membrane. However, the accurate automatic segmentation of complex images is a crucial step that has not yet been completely achieved. For this reason, 3D neuron reconstruction is generally based on interactive tracing of the neuron elements from a stack of microscopy images. This task is typically carried out by an operator who marks control points along neurites (neuron dendrites and axon) [Zhao et al. 2011], [Long et al. 2012]. Figure 1(a) shows a general view of the main parts of a pyramidal neuron.

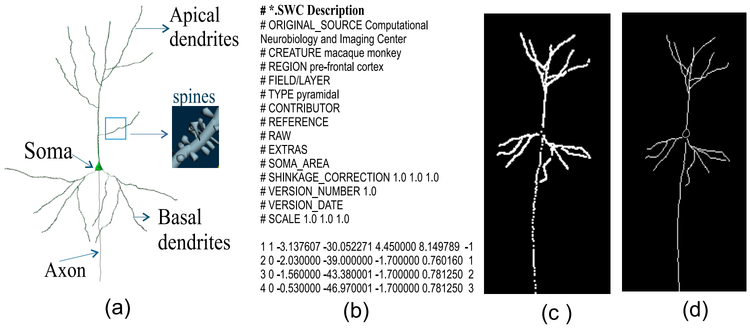

Figure 1: a) Parts of a pyramidal neuron: cell body (soma), dendrites (apical and basal dendrites) and axon. b) SWC text file describing the morphologic points (position and diameter). c) Depiction of the traced morphological points. d) Visualization of the morphological skeleton defined by the points.

The description of the cell is stored in this way as a set of 3D points, also called morphological points, which follow the trajectory of the neurites. Additionally, at each morphological point, the diameter of the corresponding neurite is stored. In the case of the soma (or cell body), the silhouette of the 2D projection is traced and stored in much the same way as dendrites. However, some laboratories only store the centre point and the diameter of a sphere that approximates the soma. The correct linking of the morphological or control points defines the morphological skeleton of the cell and constitutes the anatomical description of each segmented neuron (see figure 1(c,d)).

It should be noted that the scientific community has a huge database of neurons represented in this way. Some simulation environments also use similar neuron descriptions for creating complex neuronal circuits, which are used for large-scale simulations [Lasserre et al. 2012]. The morphology of the main elements of a neuron (axon, dendrites, spines, etc.) is essential for understanding many topological and functional aspects of neuron networks [Kuß et al. 2012]; indeed morphology plays a key role in information processing and brain function [Gleeson et al. 2007, Häusser et al. 2003]. Thus, 3D visualization represents a powerful tool that can facilitate the analysis of neuronal data [Long et al. 2012]. Previous attempts to generate 3D models to be included in standard visualization algorithms can be found in [Brito et al. 2013] and [Lasserre et al. 2012]. Nevertheless, the need to use 3D models is not always evident, given the complexity of 3D neuron models and the vast number of cells per unit of tissue volume, which results in cluttered, difficult to interpret images if rendered using conventional procedures.

2.2 3D visualization

The scientific and technological advances that are continuously taking place allow the development of new methods and techniques that contribute to the study of a wide range of systems from diverse areas. Visualization is one of the techniques which has contributed to these advances having been successfully applied to the analysis and visualization of medical data [Zhang et al. 2001], [Kuß et al. 2012], [Rautek et al. 2007], and [Dong, et al. 2003], making use of the human visual system capacity to extract information from visual scenarios. Scientific visualization and other related areas, such as visual analytics and information or data visualization [Bruckner and Gröller, 2007], are extremely active multidisciplinary research fields focused on providing graphical representations of data that increase the understanding of the phenomenon being observed [Upson et al. 1989]. The rendering technique chosen plays a crucial role in the appearance of the final images. Many factors influence the applicability of a particular strategy. For example, the nature of the input data influences, whether image-based rendering or model-based rendering is used [Kowalski et al. 1999, Kang et al. 2006]. In the latter case, the method used for representing the objects that compose the scene results in different rendering approaches such as volume, point-based or mesh renderers [Kobbelt and Botsch, 1999].

Rendering objects from compact representations, without the need of storing heavy polygonal meshes, is not a novel approach. For example, [Gumhold, 2003] explored the rendering of ellipsoids used in visualization and modelling; [Reina et al. 2005] presented a technique for the visualization of molecular dynamics; and [Stoll et al. 2005] developed a technique for the stylized rendering of line primitives. Ray-casting approaches have also been widely explored [Garrity, 1990], [Lacewell et al. 2008] and [Beyer et al. 2014]; Krüger and Wetermann were among the first to perform ray-casting on GPU [Kruger and Westermann. 2003]; Toledo *et al.* presented a study based on ray-casting of implicit surfaces on GPUs [Toledo et al. 2007], and bounding volume hierarchies (BVH) based approaches were described in [Wald et al. 2007]. The work presented in this paper is a domain-specific adaptation of ray-casting techniques that allow the skeleton-based cell descriptions obtained from neuroscience laboratories to be processed directly, avoiding the need for other standard representations commonly used in graphics.

The emergence of non-photorealistic techniques (NPR) has started a new trend in visualization styles [Gooch et al. 1998], [Gooch and Gooch, 2001], [Hertzmann, 1999], [Bram et al. 2012], and [Strothotte and Schlechtweg, 2002]. NPR methods can render images following a specific artistic style, or more in line with the present paper, can attempt to improve the clarity of the rendered images, or the way their information content is transmitted [Luft et al. 2006], [Saito and Takahashi, 1990] and [Ritschel et al. 2009]. In contrast with realistic rendering methods, NPR is not focused on reproducing the detailed appearance of objects, but on enhancing the image clarity by increasing the perceptual quality of the visualization [Sayeed and Howard, 2006].

A previous attempt to visualize brain data by applying illustrative techniques was carried out by [Choudhury et al. 2009], [Conturo et al. 1999], and [Everts et al. 2009], who suggested the application of halos for the rendering of dense line data and, more specifically, for DTI fibre tracts. The combination of colours as a crucial factor in visual aesthetics was explored by [Wang et al. 2009] via perceptual colour maps. The work described in the present paper goes a step further, not only demonstrating a domain-specific rendering for neuronal cell descriptions, but also incorporating illustrative techniques to increase the visual interpretability of complex scenes. In addition, the feasibility of parallelizing the algorithms has been assessed, opening the way towards the real time NPR rendering of complex neuronal scenes.

3 Method description

The method presented here uses a ray-casting approach, taking the morphological description of neurons as input data (as already explained in section 2.1), avoiding the need for polygonal meshes and generating a 2D rendered image that incorporates illustrative techniques in order to enhance the interpretability of the scene. Specifically, the following factors are included in the shading step to compute the final colour of a pixel:

- Tone selection according to the camera-object distance.
- Saturation adaptation according to the distance to the external contour of the object.
- Inclusion of silhouettes in the neurons that are close to the camera.
- Highlighting of relevant features such as bifurcations in dendrites and connection points between dendrites and somas.

An ad-hoc ray-casting method has been designed for computing the intersections between rays and dendrites and between rays and somas, and for assigning the proper colour to the pixels of the final image. These colours are not computed by applying standard shading techniques, but according to the illustrative strategies already mentioned. The following subsections provide a detailed explanation of the rendering algorithm and its parallel implementation.

3.1 Initial Data Representation

As previously stated, one of the contributions of this paper is the capability of rendering neurons directly from the morphological description extracted in neuroscience laboratories, which can later be included in complex circuits composed of a large number of neurons in current simulations. The input data (neurons) are not described following the common representation models in 3D graphics (polygon meshes, volumetric models, etc.). Instead, they are modelled using the morphological representations (morphological tracings describing the skeletons of cells). The morphology data of neurons is stored in *.SWC format [Cannon et al. 1998]. This SWC file describes the morphological skeleton as a set of 3D points, the diameter at each point and the connectivity between segments. The soma (or cell body) is usually described as a centre point and a diameter.

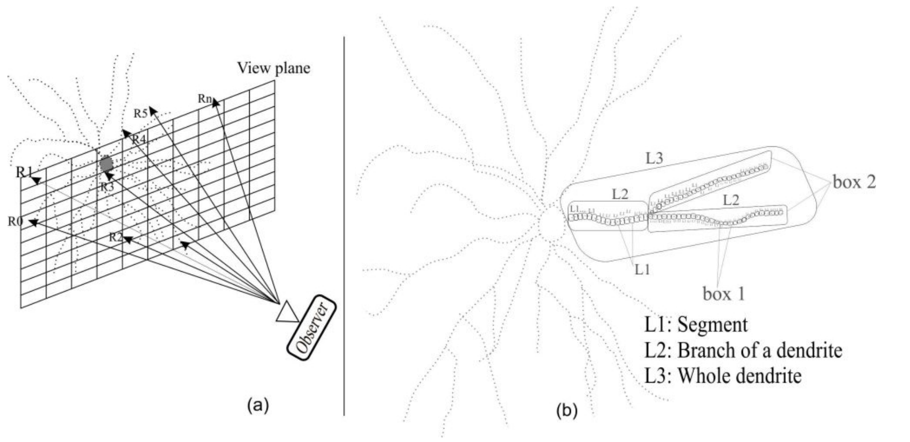

Figure 2: a) Ad-hoc ray-casting for the morphologic description of neurons. b) A bounding volume hierarchy for a neuron (BVH).

Applying the ray-casting approach involves computing intersections between the rays (shot from the camera through each pixel of the image) and each segment of the neuron or the sphere that describes its soma. With the aim of reducing the number of intersections to be checked, a bounding volume hierarchy (BVH) is computed to allow most of the segments to be ruled out when checking the intersection of a ray with some element of the scene. These ad-hoc BVHs consist of a hierarchy of boxes containing the segments that define the neurons. Figure 2(b) depicts a part of the BVH structure. In the first level (L1), there is a container box for each segment (box 1). In the second level (L2), there is a box for each section between bifurcations (box 2). The second level (L2) thus contains a branch of a dendrite, and the third level (L3) contains the whole dendrite. This hierarchy (L1, L2, L3) is stored in order. Since the structure is hierarchic, the second level (L2) contains all the first level data (L1) in the relevant section, and the third level (L3) contains all second levels (L2) in the dendrite. The purpose of this BVH technique is to avoid having to calculate all possible intersections between rays and segments, which would involve a high computational cost. In this way, several segments are packed together within a simple geometrical outline.

3.2 Ray-casting and Shading Techniques

This step involves the two main tasks of the rendering algorithm. First, the intersections between rays and objects are computed, and second, the final colour is assigned to each of the pixels of the image (using the BVH previously described). The operations for calculating the colour of pixels are the following: if the pixel is not busy (i.e., there has been no intersection between the ray passing through the pixel and any neuron segment), the assigned colour is the background colour. If the pixel is busy, the colour calculation is performed taking into account the distance to the observer in order to assign its shade. Similarly, the distance to the median axis of each dendrite is taken into account to incorporate saturation and assign silhouettes (which are applied only for ranges of distances close to the observer). Moreover, markers are incorporated in the branches of the neuron and the connections between the dendrites and soma — to act as visual resources. The sections that follow describe the shading algorithm, which distinguishes between the following cases: joints between two segments, the body cell and the special features under consideration (bifurcations and connections to the soma).

3.2.1 Segment Shading

Neuroscientists acquire the morphology of neurites (dendrites and axon) by clicking at specific points and annotating their 3D coordinates together with the neurite width at that point. The neurite shading algorithm presented here distinguishes between the regions around these morphological points (referred to as *joints*), and the regions located between consecutive, linked joints. These latter regions will be referred to as *sections*.

Regarding the rendering algorithm, intersections between rays and skeletons are computed by exploiting the BVH initially constructed. A ray is considered to intersect a segment when the minimum distance between both is below the radius of the dendrite. Since radii are measured only at the ending points of each segment, a linear

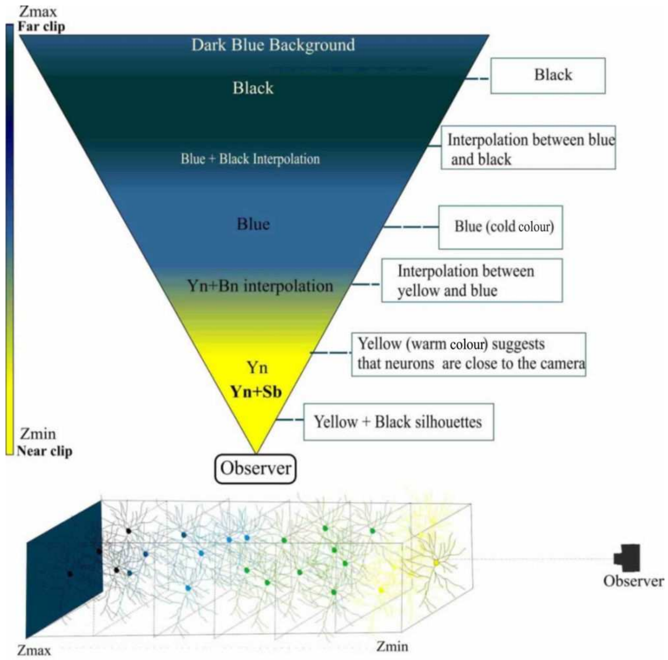

Figure 3: (Top) Hue distribution according to camera distance. (Bottom) Simple 3D scene illustrating the colours of the neurons, when the camera is located on the right-hand side.

interpolation is performed for intermediate points. Once the closest intersected segment is selected, the following step computes the pixel colour. Three factors contribute to this final colour computation: hue, saturation and the highlighting of silhouettes (whose colours depend on the object to camera distance). Figure 3 shows the hue-to-distance mapping used for this paper, where yellow has been selected for the closest objects, blue for those located in intermediate distance ranges and black for the more distant ones. Clearly, other colour mappings are possible. Within each dendrite segment, the areas near the middle axis are drawn using lower saturation levels, depending on the proximity to the dendrite's contour.

Finally, thick black silhouettes are drawn for the external contours of dendrites only in the foreground, with thickness depending on the neurite's diameter. Special attention must be devoted to the regions that are near the segments ends: this algorithm considers all of the points corresponding to the areas enclosed inside the circles in figure 4 as belonging to the two segments adjacent to them, yielding renderings that could show colour discontinuities. For this reason, the points inside the circles are considered as

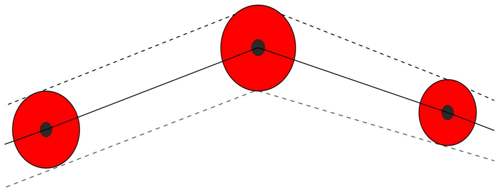

Figure 4: Regions enclosed in the red circles are considered as joints between segments.

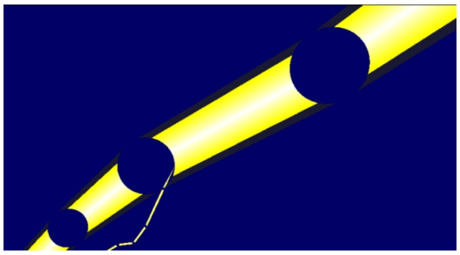

Figure 5: Segment shading. Points within each joint are shaded via interpolations between adjacent segment regions.

joints and treated in a different way. Taking this into consideration, the shading of segments produces the image shown in figure 5. The following section explains the specific algorithm used for shading these circles.

3.2.2 Shading of Joints

Even though joints are the areas corresponding to the neuroscientists' actual measurements, their much smaller size compared to the whole axon or dendrites means that it is necessary to first ensure the consistency of the representation of the segments. Later, the shading of each joint ensures smooth transitions between segments by interpolating between the pixels belonging to the adjacent segments. The red areas in figure 6 correspond to joints, and the yellow area represents the segment that links both joints. The shading of joints proceeds as follows:

Let **p** be the intersection point between a ray and the joint-sphere between two adjacent segments, described by their middle axis vectors $\overrightarrow{A}$ and $\overrightarrow{B}$ (figure 7). A new segment $\overline{pa}$ is defined, parallel to $\overrightarrow{A}$, starting at **p** and ending at **pa** (intersection between PA and the sphere). Also, $\overline{pb}$ is parallel to $\overrightarrow{B}$ and will intersect with the sphere at **pb**. The

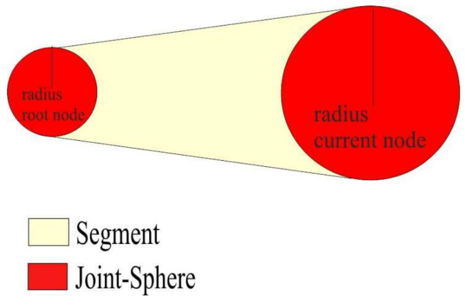

Figure 6: Joint shading. Red: regions assigned to the joints. Yellow: the region assigned to the segment.

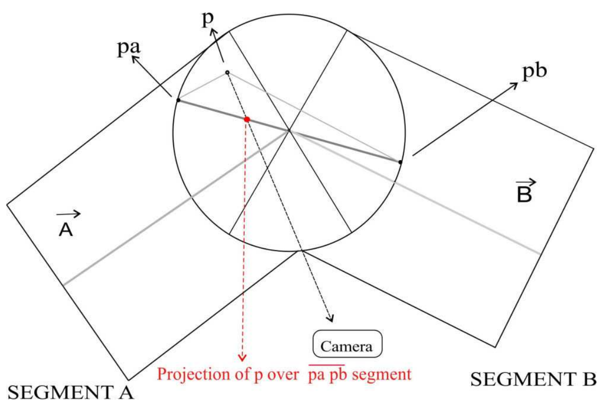

Figure 7: *Computation of the interpolation factor in joints.*

segment connecting **pa** and **pb** is used to interpolate the final colour assigned to **p**. The interpolation factor is given by the projection of **p** over this segment. Figure 8 shows the final rendering of a portion of dendrite, including the shading of segments and joints (depending on the distance to the camera).

3.2.3 Bifurcations, Soma and Connections to Soma

Somas (cell bodies) are rendered as spheres with a colour that depends on the soma-camera distance, following the algorithm explained in section 3.2.1. Bifurcations are relevant morphological features, but the extremely intricate geometry of neurites does

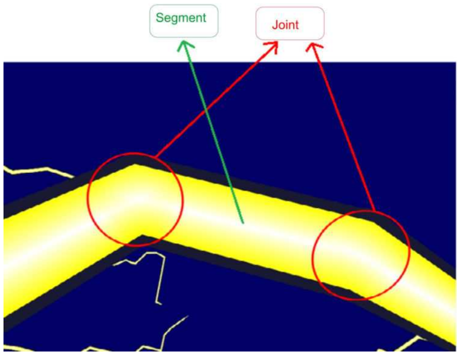

Figure 8: *Final rendering of a dendrite. Segment and joint regions are highlighted.*

not always allow them to be clearly differentiated. To alleviate this problem, a red sphere is placed over fork points, making identification of the bifurcations much easier.

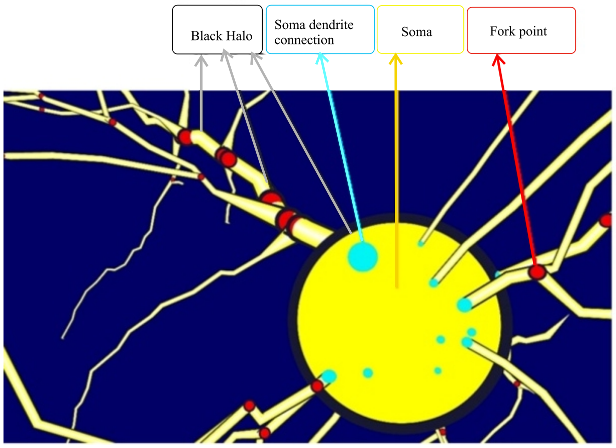

Figure 9: Enhancement of bifurcations (red circles are fork points) and connections between dendrites and soma (blue circles).

Similarly, light blue spheres can also be placed on the connections between dendrites and soma, helping users to get a better idea of the 3D location of the dendrite insertion point. Figure 9 shows a rendered soma, including all of these features.

3.3 Parallel implementation on CPU/OpenMP and GPU/CUDA

Ray-casting approaches are usually easy to parallelize. Although optimising the proposed techniques is not one of the main goals of this work, the feasibility of running them using parallel architectures opens up the possibility of achieving interactive rendering rates. Two parallel implementations have been developed: a CPU-based parallelization using OpenMP, and a GPU-based parallelization using CUDA (Compute Unified Device Architecture).

The parallelization strategy selected in this paper for the GPU implementation is based on a thread per pixel distribution, where each CUDA thread computes the final colour of a pixel in the 2D rendered image. The number of CUDA threads per block was set to 256 in order to optimize multiprocessor performance. The CUDA kernel is divided into three kernels: The first kernel computes the rays from the camera that pass through each of the pixels of the final image. The second kernel computes the intersections between each ray and the neurons of the scene. The third kernel performs the shading of the pixels and assignment of colour to each pixel. A compact representation of a neural scenario with an improved depth perception and without polygon meshes is obtained. The ad-hoc ray-casting method is easily parallelizable because every thread in CUDA performs operations for each pixel of the image. Interoperability between CUDA and OpenGL has been achieved through the use of PBOs (Pixel Buffer Objects) that are updated by the kernel and displayed afterwards through the OpenGL API. The GPU was programmed using the CUDA programming toolkit, version 4.2 on the CPU side; we used the Microsoft visual studio 2010 9.0 C++ compiler.

The hardware framework used here is based on an Intel Xeon E5645 2.40GHz with NVIDIA GeForce GTX680 (containing 8 SMXs, 1536 cores), under Windows OS. In the case of the CPU based parallelization, the OpenMP directives were inserted in the code to parallelize the computation of the final colour of pixels, including antialiasing, as well as the precomputation of the BVH (the data structure has been adapted to the morphological description of the neurons) in order to accelerate intersection tests by reducing the number of intersection tests required and also to significantly optimize the computation time (the source code is available at $http://www.gmrv.es/\sim amendoza/p01/$).

4 Results

This section presents the experimental results obtained applying the techniques described above, including antialiasing through supersampling. In addition, the results of a user preference study are also presented and discussed.

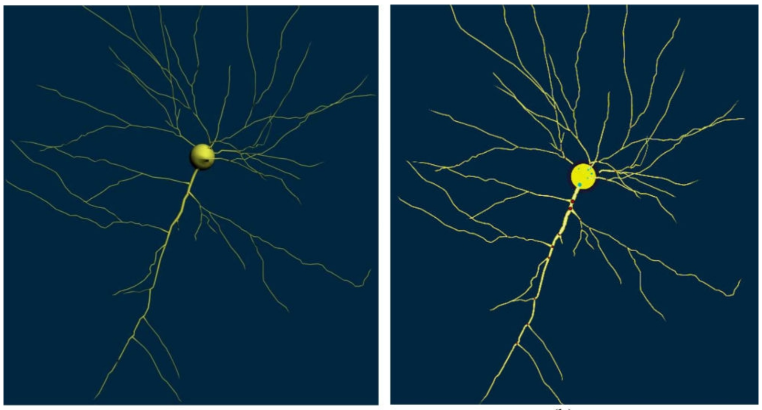

(a) (b)

Figure 10: (a) 3D visualization of a neuronal polygonal mesh applying Phong shading using commercial software. (b) NPR rendering obtained by applying the proposed techniques to the neuron's morphological description.

4.1 Experimental results

The neuron morphologies used for generating the images were downloaded from the Neuromorpho repository [Halavi et al. 2008] and were randomly selected, combining different laboratories, species and types of neurons. Therefore, the images do not faithfully reproduce a real neuronal forest in terms of the number of cells, their position or type. However, this fact does not interfere with the main goal of this work, which is to render the neurons from their original morphological descriptions incorporating illustrative techniques. Figure 8 shows a close view from a dendrite fragment. Variations in illumination are smooth enough to make the transitions between all joint-regions and segment-regions unnoticeable. Figure 10 shows a comparison between a standard rendering of a single neuron from a 3D polygonal mesh, using 3D commercial software (figure 10(a)), and the non-photorealistic rendering from the morphological description applying our method (figure 10(b)).

Highlighting the presence of bifurcations allows their position to be identified even from points of view that would otherwise result in the fork points being hidden from view; figure 11 shows an example of this. Similarly, the connections between the somas and dendrites may go unnoticed or may not be visible at all. Placing markers in these positions facilitates their identification, providing information about the diameter of the dendrites, which is also an important anatomical parameter. For comparison purposes, figure 12 shows a general view of a more complex neuronal scene. Visually following the trajectories of the dendritic and axonal trees of a neuron becomes an almost impos-

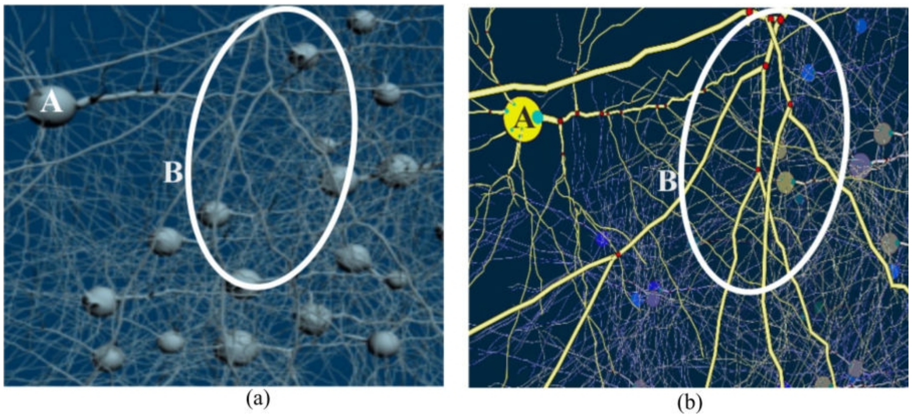

Figure 11: (a) Portion of a neuronal scene applying Phong shading to polygonal meshes. (b) Similar scene rendered from morphological skeletons, applying NPR techniques. Connections between dendrites and the soma (label A) can be clearly discerned in panel (b) while in panel 11 (a) they remain almost imperceptible. Similarly, bifurcations of the dendrite labelled as B are clearly distinguishable in panel 11(b).

sible task when analysing figure 12(a). Figure 12(b) shows an analogous scene rendered applying tone variations according to the distance to the camera. In this example, the dendrites of the selected neurons (labelled as A and B) can be easily identified, since their starting points and the cell body are at similar distances to the camera, and are depicted in a similar tone.

Moreover, the ranges of tones can be adapted to the scenario and to the region of interest in order to increase image clarity. Additionally, the presence of silhouettes seems to be very helpful to distinguish between positions when dendrites are at a similar distance from the camera, since, in this case, their positions cannot be definitively determined from tone variations alone (figure 13). In figure 13(a), it is not clear which dendrite is in front and which is behind without silhouettes, at the locations highlighted with the white ellipse. However, the presence of silhouettes in figure 13(b) clarifies the relative position of the dendrites.

Regarding computational costs, the savings in memory required to store the neuron models depend on the resolution of the polygonal meshes and the cell complexity. However, the number of triangles could vary between thousands and hundreds of thousands per neuron for intermediate resolutions, while the number of morphological points could be in the order of hundreds. Additionally, figure 14 shows a comparison between the number of polygons required for polygonal mesh representations and the

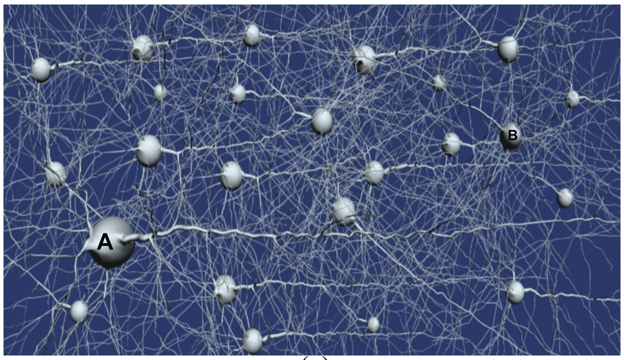

(a)

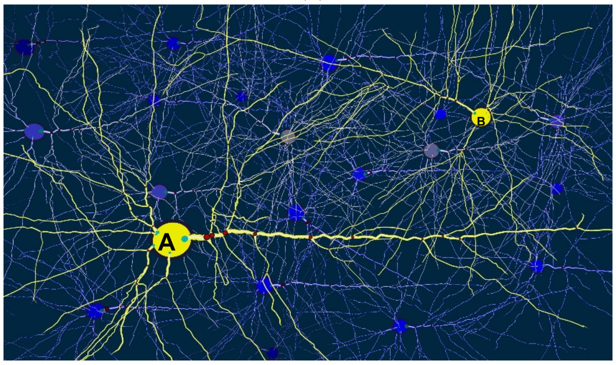

(b)

Figure 12: (a) Neuronal scene of 100 neurons (2.791.174 polygons), rendered using Phong shading. (b) Similar scene rendered by applying NPR techniques (114.907 segments). The trajectories of dendrites of neurons A and B are clearly perceivable in figure 12(b) while following them in figure 12(a) is an almost impossible task. Variation in tones according to distance makes it possible to visually declutter the image and even disregard objects that are far away.

number of segments used in the morphological descriptions, using three scenarios.

Although obtaining interactive rendering rates was not among the main goals of this paper (and consequently the methods have not been thoroughly optimized), the feasibility of parallel implementations has been explored in some depth. In order to analyse performance, two aspects have been taken into consideration, given the fact that an image-based parallelization strategy has been followed:

– Sheer geometric complexity, measured as the number of segments within the scenario.

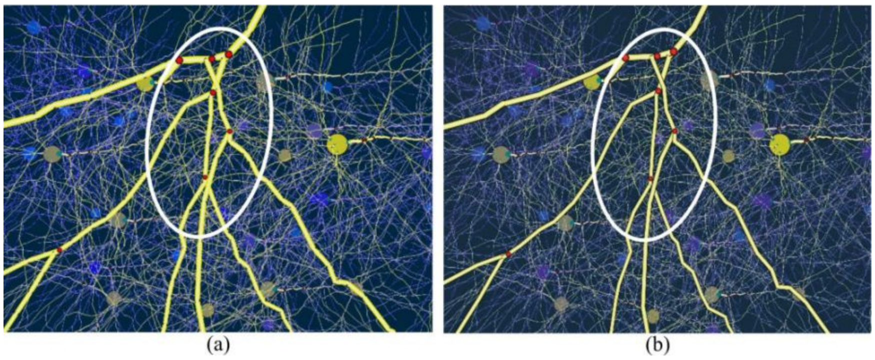

Figure 13: NPR rendering of the same scene without silhouettes (a) and with silhouettes (b). Relative positions of dendrites in cross-overs can be difficult to perceive in figure (a) (region of interest enclosed in a white circle). The inclusion of silhouettes makes it possible to clearly distinguish which dendrite is in front and which is behind.

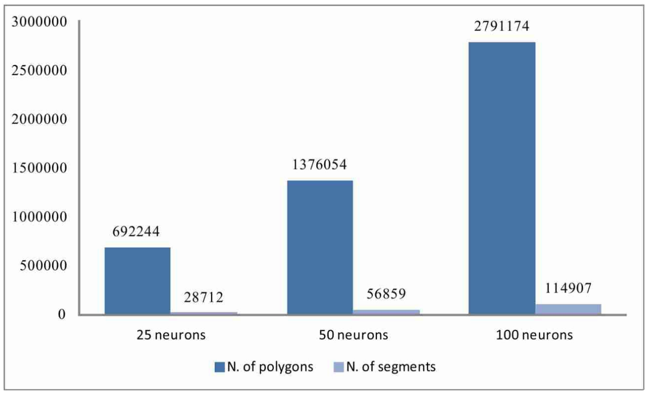

Figure 14: Comparison between the number of polygons and the number of segments in three different scenarios (25, 50 and 100 neurons).

- Workload distribution, since on the one hand there are a large number of empty pixels and on the other hand there are several other pixels that require a large number of intersection tests, this uneven distribution can negatively affect the results of a specific parallel implementation.

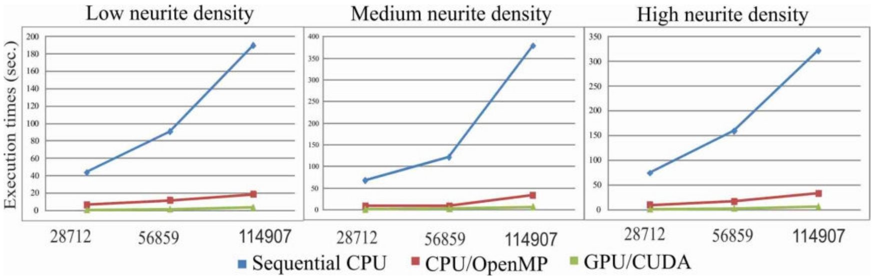

Figure 15: Evolution of execution times for different numbers of segments for 3 different implementations: sequential implementation on CPU; parallel implementation on CPU using OpenMP; and parallel implementation on GPU using CUDA. Measurements have been taken for 3 degrees of data (neurite) density.

Low density

Exec.Times (sec.) / N. of segments	Sequential	CPU OpenMP	GPU/CUDA
28712	43,65	6,84	0,82
56859	90,66	11,52	1,66
114907	189,59	18,65	3,44

Medium density

Exec.Times (sec.) / N. of segments	Sequential	CPU OpenMP	GPU/CUDA
28712	67,61	8,40	1,17
56859	121,23	8,60	2,10
114907	378,91	33,82	6,10

High density

Exec.Times (sec.) / N. of segments	Sequential	CPU OpenMP	GPU/CUDA
28712	74,08	9,47	1,24
56859	159,08	17,19	2,79
114907	321,65	33,27	6,67

Figure 16: *Execution times associated with the graphics displayed in figure 15.*

For these two reasons, the tests performed to assess the parallel implementation results include three different scenarios with a different number of segments to be rendered, but also a further three more options for the concentration of the segments, giving a total of nine test scenarios.

The experiments were performed on an Intel Xeon E645 2.40GHz (2 processors with 6 cores each with Hyperthreading with 24 cores) with Nvidia Geforce GTX 680 GPU. The CPU had 48GB of DDR2 RAM. The execution times do not include the time for reading the input files describing the scenario. Figure 15 shows a comparison of the execution times for the three different implementations: sequential; parallel using OpenMp; and parallel using CUDA (see figure 15 and figure 16). It can be observed that performance is significantly improved for both parallel implementations. The GPU version provides the best execution times, although the parallelization using OpenMP could be an interesting alternative for CPU based architectures. In terms of scalability, the performance gap becomes more evident as the complexity of the scene increases.

4.2 Preference study

A user preference study was conducted to obtain a subjective evaluation with the neural scenarios. The proposed method, based on NPR techniques, was compared with realistic techniques. The user preference study was performed focusing on the following aspects:

- The distance perception from closer neurons (foreground) to further neurons (background).
- Cell morphology and special neuron features, such as branches; insertion points of dendrites into the soma; fork points, etc.
- Occlusions, specifically when dendrites cross over other neurons' neurites.

4.2.1 Experimental design

A set of neurons of different types was selected from the NeuroMorpho repository, and six scenarios were rendered with these data (*.swc). Three of them were rendered with a 3D commercial software package, using 3D polygonal models generated with the Neuronize tool [Brito et al. 2013] applying Phong shading. The three other scenarios were rendered with the present illustrative techniques. Six questions were defined for the study; each question had two similar scenarios, one rendered with the proposed NPR techniques (scenario 2) and the other rendered applying the standard Phong shading model (scenario 1). In each question, each participant had to indicate their degree of agreement or disagreement with the comparisons made, using a Likert scale (1-5). The questions were:

1. Observe the following scenarios. In order to identify the four closest neurons to the camera, scenario (2) is more useful than scenario (1).
2. In order to identify the dendrites leading from the marked soma, scenario (2) is more useful than scenario (1).
3. Observe the dendritic branching path labelled within the marker. Scenario (2) makes it easier to distinguish the branches and trajectory than scenario (1).

4. The bifurcations of the dendritic tree are easier to distinguish in scenario (2).

5. Observe the dendritic tree enclosed in a white circle. The crossings between dendrites are easier to perceive in scenario (2).

6. Observe the dendritic tree enclosed in the white circle. The dendritic silhouettes are easier to perceive in scenario (2).

4.2.1.1 Populations and procedure

To validate the method, several participants (N=25) were invited to perform a subjective study of neuronal scenarios. 70% of the participants had a computer graphics background, and 2 participants were physicians. Each participant answered six subjective questions in which two scenarios were compared. Each of the questions involved a scenario rendered using realistic techniques, and also using the methods presented here. In addition, a statistical analysis with SPSS 17.0 was performed; specifically, the chi-square test was used to evaluate the significant differences between the users' preferences.

4.2.1.2 Results and discussions

The study results showed that the present method is significantly better than the realistic techniques. Over 90% of the participants agreed that the scenarios with illustration techniques improved their perception in terms of the visualization of branches of neurons and morphological features. Figure 17 shows the study results, where a significant difference between how the users prefer the visualizations with NPR-rendered scenarios over those rendered with standard techniques was found.

In the first question, 96% of the participants strongly agreed or agreed that scenario (2) clearly improved visualization compared to scenario (1). However, 4% of the participants agreed that scenario (1) was better to distinguish the neurons closer to the observer. The statistical results, $\chi^2 = 18,32, df = 2, p < 0.001$, showed a significant difference between the users' preferences. On the second question, 88% of the participants strongly agreed or agreed that with scenario (2) it is possible to observe the dendrites leading from the soma better, while 8% were neutral, and 4% disagreed with the statement and preferred scenario (1). The statistical results, $\chi^2 = 17,40, df = 3, p = 0.001$, showed a significant difference between the users' preferences. In the third question, 100% of the participants strongly agreed or agreed that the dendritic branching pattern in scenario (2) is more easily distinguished than in scenario 1. The statistical results, $\chi^2 = 11,56, df = 1, p = 0.001$, showed a significant difference between the users' preferences. Similarly, in the fourth question, 100% of participants totally agreed or agreed that in the scenario (2) bifurcations of the dendritic tree are easier to distinguish than in scenario (1). The statistical results, $\chi^2 = 14,44, df = 1, p = 0.001$, showed a significant difference between the users' preferences. In the fifth question, 80% of the participants strongly agreed or agreed that in scenario (2) they could perceive the crossings between dendrites better than with scenario (1). 16% of the participants were neutral, while 4% disagreed with the statement. The statistical results, $\chi^2 = 11.00, df = 3, p = 0.012$, did not show a significant difference between the users' preferences.

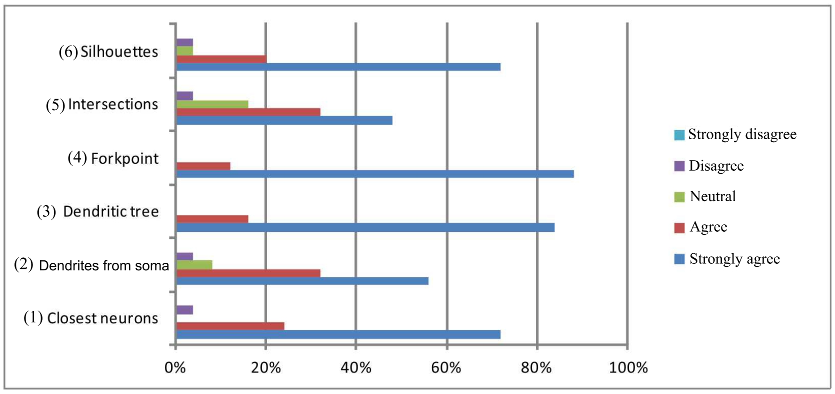

Figure 17: Plot of the user preferences. We found a significant difference between how the users prefer the visualization with scenarios with NPR techniques over the scenarios with the realistic techniques.

Finally, in the sixth question, 92% of the participants strongly agreed or agreed that they could observe the dendritic shapes better with scenario (2) than with scenario (1). 4% of the participants were neutral, while the remaining 4% disagreed with the statement. The statistical results, $\chi^2 = 31, 16, df = 3, p = 0.001$, showed a significant difference between the users' preferences.

In general, the user study shows quite large differences in user preferences. In particular, in figure 17 a clear tendency of preferring the scenarios rendered with the method presented here can be observed. These results suggest that there are significant differences in user preferences between the scenarios with illustrative technique and realistic techniques. More than 85% of the participants preferred the NPR technique because it improves the visualization and interpretation in complex scenarios, especially in occlusion situations among dendrites and neurons where they have to discern parts of neurons such as dendrite branches with their fork-point features.

5 Conclusions and Future Work

The work presented here includes progress made on two important fronts: First, the proposed rendering technique takes the morphological descriptions commonly extracted by the neuroscience laboratories as input, thereby avoiding the need to generate and store other intermediate representations such as polygonal meshes, specifically oriented towards 3D visualization algorithms. This feature is especially significant when considering the burden of visualizing even small fractions of the huge number of neurons present in the brain. The second main contribution of this work is the incorporation of illustrative techniques in the implemented rendering algorithms. This allows the generation of enriched images that facilitate the analysis and understanding of complex scenes or intricate circuits of neurons. A simple visual inspection of the results shows that depth perception is better when applying tone variations according to the distance to the camera. The inclusion of thick silhouettes for neurons closer to the camera seems to help in identifying neurite trajectories and in distinguishing relative positions, particularly whenever they cross each other. Clearly, highlighting certain features —such as bifurcations or connections between neurites and soma— on neuronal scenarios can make them more discernible, thereby improving the overall visualization.

Although a rigorous user performance validation has not yet been carried out, an initial subjective user preference study has already been conducted. The results show that there is a significant difference in the users' preference for the NPR visualizations over the realistic renderings. The approach presented here should be considered as an alternative visualization framework that could be combined with other rendering strategies that may be more suitable for certain purposes. Combining different rendering techniques together into a multi-view visualization framework will allow the most appropriate representation to be selected according to the user needs. The authors are also aware of the need for 3D surface representations in some cases; nevertheless, the proposed technique is a novel rendering approach in neuroscience that warrants further investigation. The work presented here may be extended in the following ways:

- Performing a rigorous evaluation to study the impact of NPR techniques on user performance for a set of visual inspection tasks.
- Analysing other illustrative techniques that facilitate the analysis and interpretation of complex neural scenes.
- Optimizing the implemented algorithms in order to achieve interactive visualizations. In this regard, GPU shaders and other parallelization strategies should also be explored.
- Integrating the proposed technique into a standard visualizer, thereby making it possible to combine different visualization approaches.

Acknowledgements

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project), the Spanish Ministry of Economy and Competitiveness (Cajal Blue Brain Project, Spanish partner of the Blue Brain Project initiative from EPFL and grant (TIN2014-57481). Also we thank Juan Ramón Cañizares and Angel Merino for their helpful insight and collaboration with the neural 3D scenarios.

References

[Bram et al. 2012] Brambilla, A. and Carnecky, R. and Peikert, and R. and Viola, I. and Hauser, H.: "Illustrative Flow Visualization: State of the Art, Trends and Challenges"; Eurographics STAR Reports, (2012), 75-94.

[Brito et al. 2013] Brito, J. and Mata, S. and Bayona, S. and Pastor, L. and DeFelipe, J. and Benavides Piccione, R.: "Neuronize: A tool for building realistic neuronal cell morphologies"; Frontiers in Neuroanatomy 7(15) (2013), ISSN 1662-5129.

[Bruckner and Gröller, 2007] Bruckner, S., and Gröller, M. E.: "Style transfer functions for illustrative volume rendering"; Computer Graphics Forum, 26(3), (2007), 715-724.

[Cannon et al. 1998] Cannon, R. C., and Turner, D. A., and Pyapali, G. K., and Wheal, H. V.: "An on-line archive of reconstructed hippocampal neurons"; Journal of Neuroscience Methods, 84(1–2), (1998), 49-54.

[CarlZeiss, 2014] "Microscopy, Carl Zeiss Advanced Imaging" [Online]. LSM510 Meta-Laser Scanning Microscope, URL: $http : //www.zeiss.com/microscopy/en_de/products/$ (2014).

[Casares and Gnauck, 2009] Casares, A., and Gnauck, P.: "CrossBeam, principles and application"; Proceedings of the 2009 Third International Conference on Quantum, Nano and Micro Technologies. February 1-7, Publishing, Cancun, Mexico, (2009), 122-124.

[Choudhury et al. 2009] Choudhury, A. and N. M. I., and Parker, S. G.: "Ray tracing NPR-style feature lines"; the NPAR '09: Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering, New Orleans, Louisiana, (2009), 5-14.

[Conturo et al. 1999] Conturo, T. E., and Lori, N. F., and Cull, T. S., and Akbudak, E., and Snyder, A. Z., and Shimony J.S: " Tracking neuronal fiber pathways in the living human brain"; Neurobiology, Applied Physical Sciences, (1999), 10422-10427.

[CUDA, 2015] Nvidia: "Home page maintained by nvidia", URL:$http : //www.nvidia.com/$ (Accessed March 2015).

[Debevec, 2008] Debevec, P.: "Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography"; ACM SIGGRAPH, Los Angeles, California. (2008), pp. 32:1-32:10.

[Dong, et al. 2003] Dong, F., Clapworthy, G. J., Lin, H., and Krokos, M. A.: "Nonphotorealistic rendering of medical volume data"; IEEE Comput.Graph.Appl., 23(4), (2003), 44-52.

[Everts et al. 2009] Everts, M. H., and Bekker, H., and Roerdink, J. B. T. M., and Isenberg, T.: "Depth-dependent halos: Illustrative rendering of dense line data"; IEEE Computer Society, 15(6), (2009), 1299-1306.

[Garrity, 1990] Garrity, Michael P.: "Raytracing irregular volume data". In Proceedings of the 1990 workshop on Volume visualization (VVS '90); ACM, New York, NY, USA, (1990), 35-40.

[Gleeson et al. 2007] Gleeson, P., and Steuber, V., and Silver, R. A.: "NeuroConstruct: A tool for modeling networks of neurons in 3D space"; Neuron, 54(2), (2007) 219-235.

[Gooch and Gooch, 2001] Gooch, B. and Gooch, A.: "Non-Photorealistic Rendering", A K Peters CRC Press. ISBN-13: 978-1568811338 (2001), 254p.

[Gooch et al. 1998] Gooch, A., and Gooch, B., and Shirley, P., and Cohen, E.: “A Non-photorealistic Lighting Model for Automatic Technical Illustration”; Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, (1998), 447-452.

[Gumhold, 2003] Gumhold, S.: “Splatting Illuminated Ellipsoids with Depth Correction”; Proc. Vision, Modeling, and Visualization (2003), 245-252.

[Häusser et al. 2003] Häusser, M., and Mel, B.: “Dendrites: Bug or feature?”; Current Opinion in Neurobiology, 13(3), (2003), 372-383.

[Halavi et al. 2008] Halavi, M., and Polavaram, S., and Donohue, D. E., and Hamilton, G., and Hoyt, J., Smith, and K. P., et al.: “NeuroMorpho.org implementation of digital neuroscience: Dense coverage and integration with the NIF”; Neuroinformatics, (2008), 241-252.

[Hermosilla et al. 2010] Hermosilla, P., and Brecheisen, R., and Vázquez, P., and Vilanova, A.: “Uncertainty Visualization of Brain Fibers”; Proc. CEIG Publishing, Spain, (2010) p.31-40.

[Hertzmann, 1999] Hertzmann, A.: “Introduction to 3D Non-Photorealistic Rendering: Silhouettes and Outlines”; Non-Photorealistic Rendering, SIGGRAPH Course Notes, (1999).

[Beyer et al. 2014] Beyer, J. and Hadwiger, M. and Pfister, H. : “A Survey of GPU-Based Large-Scale Volume Visualization,” Eurographics Conference on Visualization (EuroVis). (2014).

[Kang et al. 2006] Kang, S. B., and Li, Y., and Tong, X., and Shum, H.: “Image-based rendering. Found”;Trends.Comput.Graph.Vis., 2(3), (2006), 173-258.

[Kobbelt and Botsch, 1999] Kobbelt, L., and Botsch, M.: “A survey of point-based techniques in computer graphics”; Computers & Graphics, 28, 801-814 (1999).

[Kowalski et al. 1999] Kowalski, M., and Markosian, L., and Northrup, J., and Bourdev, L., and Barzel, R., and Holden, L., et al.: “Art-based rendering of fur, grass, and trees”; Procedings of ACM SIGGRAPH 99, ACM Press / ACM SIGGRAPH, (1999).

[Kruger and Westermann. 2003] Kruger, J. and Westermann, R.: “Acceleration Techniques for GPU-based Volume Rendering”. In Proceedings of the 14th IEEE Visualization 2003 (VIS’03) (VIS ’03); IEEE Computer Society, Washington, DC, USA, (2003), 38-.

[Kuß et al. 2012] Kuß, A., and Preim, B.: “Design and analysis of visualization and browsing methods for spatial neuroanatomical atlases”; Journal 1. Doktorandentagung Magdeburger-Informatik-Tage, MIT, 27, (2012).

[Lacewell et al. 2008] Lacewell, D. and Burley, B. and Boulos, S and Shirley, P.: “Raytracing prefiltered occlusion for aggregate geometry”; Proceedings of the IEEE Symposium on Interactive Raytracing, (2008).

[Laine et al. 2010] Laine S., and Karras, T.: “Two methods for fast ray-cast ambient occlusion”; Computer Graphics Forum (Proc.Eurographics Symposium on Rendering 29(4), (2010).

[Lasserre et al. 2012] Lasserre, S., and Hernando, J., and Hill, S., and Schuermann, F., and de Miguel Anasagasti, P., and Jaoude, G. A., et al.: “A neuron membrane mesh representation for visualization of electrophysiological simulations. Visualization and Computer Graphics”; IEEE Transactions on, 18(2), (2012), 214-227.

[Lehtinen et al. 2012] Lehtinen, J., Aila, T., Laine, S., & Durand, F..: “Reconstructing the indirect light field for global illumination”; ACM Transactions on Graphics, 31(4) (2012).

[Long et al. 2012] Long, F., Zhou, J., and Peng, H.: “Visualization and Analysis of 3D Microscopic Images”; PLoS Computational Biology 8, (2012).

[Luft et al. 2006] Luft, T., and Colditz, C., and Deussen, O.: “Image enhancement by unsharp masking the depth buffer”; ACM Trans.Graph., 25(3), (2006), 1206-1213.

[Markram, 2006] Markram H.: “The Blue Brain Project - BBP”; Nature Rev. Neuroscience. Feb;7(2), (2006), 153-160.

[Mcardle, 2014] McArdle, O.: “The Human brain project - HBP” [*Online*]; URL: *https://www.humanbrainproject.eu/* (2015).

[Nichols and Wyman, 2010] Nichols, G., and Wyman, C.: “Multiresolution splatting for indirect illumination”; Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games, Boston, Massachusetts, (2009), 83-90.

[Otten et al. 2010] Otten, R., and Vilanova, A., and van de Wetering, H.: “Illustrative white matter fiber bundles”; Comput. Graph. Forum, 29(3), (2010), 1013-1022.

[Rautek et al. 2007] Rautek, P., Bruckner, S., and Groller, E.: "Semantic layers for illustrative volume rendering. Visualization and Computer Graphics", IEEE Transactions on, 13(6), (2007), 1336-1343.

[Reina et al. 2005] Reina, G., and Ertl, T.: "Hardware-accelerated glyphs for mono- and dipoles in molecular dynamics visualization"; Proceedings of the Seventh Joint Eurographics / IEEE VGTC conference on Visualization, Leeds, United Kingdom. (2005), 177-182.

[Ritschel et al. 2009] Ihrke, M., and Ritschel, T., and Smith, K., and Grosch, T., and Myszkowski, and K., and Seidel, H.: "A Perceptual Evaluation of 3D Unsharp Masking"; Human Vision and Electronic Imaging XIII. Volume 7240 of SPIE Proceedings, (2009).

[Roettger et al. 2003] Roettger, S. and Guthe, S. and Weiskopf, D. and Ertl, T. and Strasser, W.: "Smart hardware-accelerated volume rendering"; In Proceedings of the symposium on Data visualisation (VISSYM '03). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, (2003), 231-238.

[Toledo et al. 2007] Toledo, R.D. and Lévy, B. and Paul, J. : "Iterative Methods for Visualization of Implicit Surfaces on GPU."; In ISVC, International Symposium on Visual Computing - Lecture Notes in Computer Science, Anonymous (2007).

[Ruiz et al. 2009] Ruiz, A. and Ujaldón, M. and Cooper, Lee and Huang, K.: "Non-rigid Registration for Large Sets of Microscopic Images on Graphics Processors"; J. Signal Process. Syst. 55, 1-3 (April 2009), (2009), 229-250.

[Saito and Takahashi, 1990] Saito, T., and Takahashi, T.; "Comprehensible rendering of 3-D shapes"; SIGGRAPH Comput. Graph., 24(4), (1990), 197-206.

[Sayeed and Howard, 2006] Sayeed, R. and Howard, T. "State of the Art Non-Photorealistic Rendering (NPR) Techniques"; Theory and Practice of Computer Graphics 2006, (2006), 89p..

[Stoll et al. 2005] Stoll, C., Gumhold, S., and Seidel, H.: "Visualization with stylized line primitives"; IEEE Visualization, (2005), 88p.

[Strothotte and Schlechtweg, 2002] Strothotte, T. and Schlechtweg, S.: "Non-Photorealistic Computer Graphics: Modeling"; Rendering, and Animation. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. ISBN 1-55860-787-0, (2002).

[Upson et al. 1989] Upson, C., and Faulhaber, T. and A., Jr., and Kamins, D., and Laidlaw, D., and Schlegel, D., and Vroom, J., et al. "The Application Visualization System: A Computational Environment for Scientific Visualization"; Computer Graphics and Applications, IEEE, 9(4), (1989), 30.

[Wald et al. 2007] Wald, I., and Mark, W. R., and Gunther, J., and Boulos, S., and Ize, T., and Hunt, W., et al.: "State of the art in ray tracing animated scenes"; STAR Proceedings of Eurographics 2007, 28(6), (2007), 1691-1722.

[Wang et al. 2009] Wang, L., and Giesen, J., and McDonnell, K. T., and Zolliker, P., and Mueller, K.: "Color design for illustrative visualization"; IEEE Transactions on Visualization and Computer Graphics, 14(6), (2009), 1739-1754.

[Wolf, 2013] Wolf L.: "White house unveils $100 million BRAIN initiative". Chemical & Engineering News. ISSN 0009-2347, (2013).

[Zhang et al. 2001] Zhang, S., and Demiralp, C., and Keefe, D. F., and Basser, P. J., and Pierpaoli, C., and Chiocca, E. A.: "An Immersive virtual Environment for DT-MRI Volume Visualization Applications: a Case Study; In Proceedings of the conference on Visualization'01", Apr., Anonymous IEEE Computer Society, San Diego, California, (2001), 437-440.

[Zhao et al. 2011] Zhao, T., Xie, J., Amat, F., Clack, N., Ahammad, P., Peng, H., Long, F. and Myers, E.: "Automated Reconstruction of Neuronal Morphology Based on Local Geometrical and Global Structural Models"; Neuroinformatics, 9 (2-3), issn 1539-2791 (2011), 247-261.