

A Neuron Membrane Mesh Representation for Visualization of Electrophysiological Simulations

Sébastien Lasserre, Juan Hernando, Sean Hill, Felix Schürmann, Pedro de Miguel Anasagasti, Georges Abou Jaoudé, and Henry Markram

Abstract—We present a process to automatically generate three-dimensional mesh representations of the complex, arborized cell membrane surface of cortical neurons (the principal information processing cells of the brain) from nonuniform morphological measurements. Starting from manually sampled morphological points (3D points and diameters) from neurons in a brain slice preparation, we construct a polygonal mesh representation that realistically represents the continuous membrane surface, closely matching the original experimental data. A mapping between the original morphological points and the newly generated mesh enables simulations of electrophysiological activity to be visualized on this new membrane representation. We compare the new mesh representation with the state of the art and present a series of use cases and applications of this technique to visualize simulations of single neurons and networks of multiple neurons.

Index Terms—Visualization techniques and methodologies, curve, surface, solid, and object representation, data mapping, neuronal network visualization.

1 INTRODUCTION

THERE is a growing interest in biologically detailed brain simulation and visualization [2], [28], [21]. The principal information processing cell in the brain—the neuron—is characterized by remarkably complex arbors of dendrites (which receive signals from other neurons) and axons (which send signals to other neurons). Indeed, each neuron is unique, but many types of neurons can be distinguished by the distinctive shapes of their dendritic and axonal arbors. Each neuron is also an electrical entity—converting chemical synaptic signals into electrical voltage impulses that propagate along the dendritic arbors as part of a complex decision making process that determines whether a neuron relays an electrical “spike” along its axon. Therefore, a key challenge in neuroscience is to relate the structure of neurons to their electrical activity in order to understand neural function and information processing in the brain.

Measuring and reconstructing the complex morphology of an entire neuron is tedious and time consuming. The detailed shapes of neurons are too minute to be seen by the human eye and the cells themselves are translucent. When

filled with a stain to make them visible, bright field microscopy [36] can reveal the detailed anatomy of these cells at magnification factors of 10-100 \times , including the soma and the arborization of the dendrites and axons.

For over a century, the shape of neurons has been reconstructed using light microscopy. Although previously hand-drawn on paper, these reconstructions are now captured with computer-aided digitization by human operators [13]. Other imaging techniques, such as confocal [6] and 2-photon [33] microscopy, have the potential to provide further detailed insights into topological details of the neuronal membrane, but the techniques are prohibitively costly and slow, making it difficult to acquire many whole neuron samples [29]. Electron beam (EM) microscopy [11], [7], [3] is the most accurate imaging technology to date, but it operates at a smaller scale and is not currently suitable for whole cell reconstruction.

Digitized light microscopy reconstructions consist of a set of 3D points and diameters that provide a nonuniform discrete sampling of neuronal fibers. Although, they provide an approximate and nonuniform sampling of the continuous shape of the neuronal membrane, the morphological points faithfully capture the complex details of the neuron arborization. This detailed point digitization process is thus essential to visualize neuron morphologies and these reconstructions represent the most readily available digital anatomical data for whole neurons [16].

The essential anatomical features of neurons captured through morphological reconstruction are then used for detailed electrophysiological simulations of voltage dynamics throughout the 3D structure of the neuron (Neuron [5], Genesis [39], Neuroconstruct [14]). The visualization of such spatiotemporal data represents a particular challenge because of the complexity of neuron morphologies and the limitations of the morphological point representation.

- S. Lasserre, S. Hill, F. Schürmann, and H. Markram are with the EPFL - Blue Brain Project, AAB 201 - station 15, Lausanne 1015, Switzerland. E-mail: {sebastien.lasserre, sean.hill, felix.schuermann, henry.markram}@epfl.ch.
- J. Hernando and P. de Miguel Anasagasti are with the CesViMa, Universidad Politécnica de Madrid, Spain. E-mail: {jhernando@laurel.datsi.fi.upm.es, pmiguel@fi.upm.es.
- G.A. Jaoudé is with the EPFL, Laboratoire d'Informatique et de Visualization, Bat. BP - 3ème étage, station 15, Lausanne 1015, Switzerland. E-mail: georges.abou-jaoude@epfl.ch.

Manuscript received 3 Feb. 2010; revised 3 Dec. 2010; accepted 13 Dec. 2010; published online 2 Mar. 2011.

Recommended for acceptance by P. Rheingans.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number TVCG-2010-02-0031. Digital Object Identifier no. 10.1109/TVCG.2011.55.

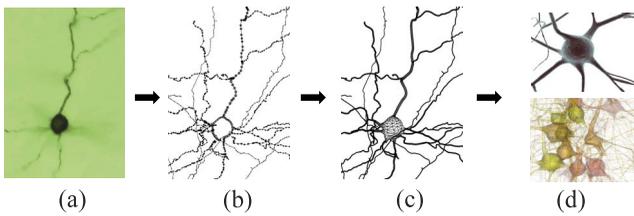


Fig. 1. Process for 3D continuous neuronal membrane generation. In (a) the bright field microscopy reveals the detailed anatomy of a stained neuron. (b) shows the morphological skeleton as captured with computer-aided digitization by human operators. Starting with this model of the morphology, we propose an extrusion-based algorithm to generate a 3D mesh membrane depicted in (c), and we describe how this new continuous membrane can be used to explore neural function and information processing in the networks of the brain (d).

Neurons are entities with smooth and continuous membrane surfaces that form complex arborizations and the morphological point representation does not provide a continuous surface representation. Here, we present a simple and general algorithm that generates a continuous representation of a neuronal membrane surface from a set of sampled morphological points. This approach provides a smooth, continuous, and high fidelity representation of neuron morphologies that can be used for new types of scientific visualizations (see Fig. 1). In addition, we provide a mapping method that relates mesh vertices to the original measured morphology, preserving the link between the generated mesh and the electrophysiological simulation data. The algorithms are tested for generality and validity against a database of nearly 200 biologically measured neuron morphologies. Finally, we present a series of use cases that demonstrate the applications of this new mesh representation to single cell and network simulations.

2 BACKGROUND

The raw data that form the basis of the proposed mesh generation process is obtained from biological experiments [16]. An individual neuron in a brain slice preparation is impaled with a pipette and a chemical dye (e.g., biocytin) is then diffused throughout the selected neuron to fill the extent of its dendritic and axonal arborizations. After chemical staining of the tissue using immunohistochemical techniques, the stained neuron is visible through increased optical contrast under bright field microscopy (see Fig. 2a). It is then possible to trace the 3D contour of a neuron and to digitally reconstruct the neuronal arbors by means of a software application such as NeuroLucida [13].

The digitization process is carried out manually by a human operator. The dendritic and axonal shapes are measured by the operator tracing the fibers with mouse clicks specifying morphological points (comprised of a 3D coordinates and a diameter—see Fig. 2b). The accuracy and sampling interval of the morphological points can therefore be highly nonuniform and dependent on the individual operator. On average 8-10,000 morphological points form a complete description of the neuron morphology.

We refer to this type of neuron reconstruction data as the *morphological point representation* or *morphological skeleton*. Tracing the structures in order results in a unique spanning tree—an acyclic graph—of the neuronal morphology. The

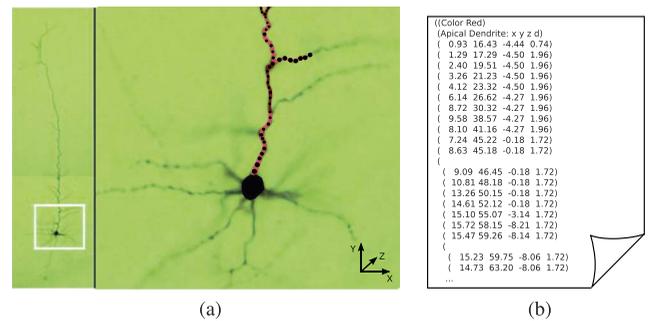


Fig. 2. 3D reconstruction of a real neuron. (a) shows bright field microscopy view of a layer VI pyramidal neuron (left) with a magnification of a portion of the neuron to show detail of the area around the cell body (soma) of the same neuron (right). Each morphological point is represented as a point with 3D location and diameter. In (b) the hierarchical data describe the morphological points of the cellular reconstruction.

root node of the graph is at the cell body (also called the soma). Two morphological points describe a conical frustum shaped segment of the neuronal morphology. The reconstructed neuronal morphology is composed of a hierarchical arrangement of these segments. An unbranched sequence of segments comprises a section of the neuron. Sections join one another at branch points. Different segments do not intersect (see Fig. 3).

As opposed to the neuronal branches, the soma is measured by a coarse outline of the maximum size (in planar projection) rather than a complete description. It is possible to draw several such “isometric” contours for the half of the soma toward the microscope to more accurately capture the shape, but due to the occlusion by the cell body itself, it is not possible to obtain the same data for the lower

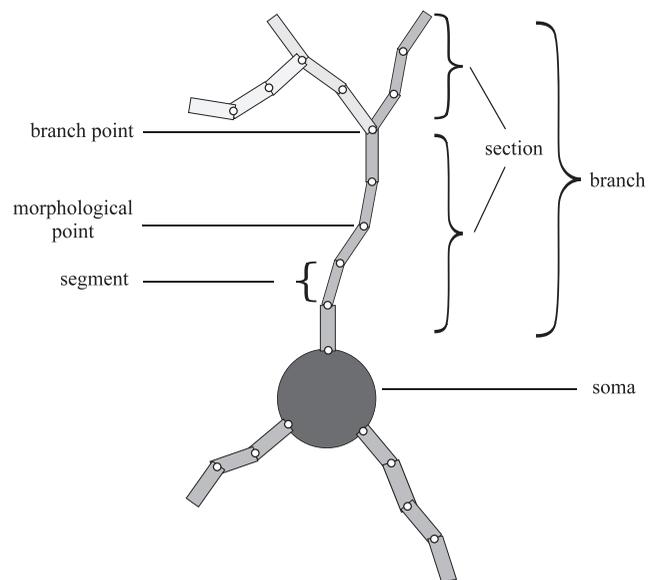


Fig. 3. Schematic view of the neuronal anatomy. 3D coordinates with an associated diameter are called morphological points. Conical frusta defined by two morphological points are called segments. A nonbranching sequence of segments defines a section while the sequence of sections following the same path delineates a branch. Child sections are sections that branch off of the parent sections at the branch point. Each branch is connected to a parent branch or, in the case of first-order branches, directly to the cell body of the neuron called soma.

part of the soma. Thus, in practice, the outline of the soma is often the only characteristic measured. Furthermore, the occlusion problem extends to the beginning of trees that originate from the soma but are occluded (in the Z-dimension). Due to the occlusion problem, it is possible that some neuronal reconstructions can contain the beginning of branches displaced in the Z-dimension.

3 RELATED WORK

The generation of an accurate closed surface from a skeleton is a common topic in computer graphics. In this section, we explore the pros and cons of three different methods to move from a skeleton representation to a surface representation: the 3D vector-fields techniques, the representation from implicit models, and the mesh extrusion techniques.

Traditionally, the 3D vector-field techniques have been used to quickly and accurately display any measurements of physical quantities using visual attributes including color, line width, texture, and orientation [32]. In the neuroscientific domain, this powerful technique is often used to visualize dense nervous tissue [19], [22], [25], but generally it merely represents white matter fiber bundles. Using these tube-based rendering techniques, the neuronal morphology can be described using minimum memory and due to the extensive progress in programmable rendering on the graphics processing unit (GPU) reasonable rendering speed is guaranteed. However, these representations do not provide specific information about the interior or the exterior of the neuron or about the three-dimensional form of the soma. In addition, the use of tube-based representations in combination with transparency induces artifacts at the junction between successive tubes.

Implicit surfaces have the potential to provide a more detailed representation of the neuronal morphology. The scalar field specifies the minimum distance to a shape, where the distance may be tested (positive or negative) to distinguish between the interior and the exterior shape. All the points at a given distance define an isosurface. Implicit models face two challenges where the faithful definition of the implicit surface constitutes the first challenge. The Adaptive Distance Field (ADF) approach can be used, for example, to benefit from an irregular sampling based on hierarchical structure or octree [12], [26] to increase the details while reducing the memory cost. The accurate sampling of these implicit models and the mesh construction constitute the second challenge. Different methods use the sampling result to render isosurfaces. The methods based on the well-known and robust Marching-Cubes algorithm [20] have to deal with overtessellated and nonwater-tight polygonal surfaces. While remeshing techniques can be used as postprocessing operations to simplify and to improve the quality of the mesh [1], alternatives like particle-based approaches [23] and advancing front techniques [30] have been developed to directly generate compact sets of nearly regular triangles and prevent the geometrically and topologically accurate output mesh from becoming unreasonably dense.

Mesh extrusion techniques to build high quality polygonal surfaces constitute another option. These techniques typically start with a coarse mesh that approximates the

object and incrementally improve the mesh geometry and triangle quality. These techniques are very flexible [38] and can produce nonuniform rational basis spline (NURBS) patches [41] or subdivision surfaces [40] using a nonshrinking algorithm [17], [34]. While NURBS patches are more efficient in terms of memory, we prefer a subdivided mesh that has been incrementally refined from a coarse mesh based on the morphological point representation because it can be rendered very efficiently taking advantage of current GPU architectures.

4 NEURON MEMBRANE MESH REPRESENTATION

The process to generate a neuron membrane mesh representation starts from the morphological point representation to extrude a coarse mesh with quadrilateral faces, which then serves as a control mesh to generate the subdivided surface that approximates the neuronal membrane surface. The algorithm proceeds in five steps:

1. Branch identification—branches are identified as sets of sections that form a continuous path;
2. Morphological point resampling—the morphological points must be resampled to adjust the density of morphological points to remove points from over-sampled portions and increase sampling near branch points to reduce smoothing artifacts as explained later;
3. Mesh extrusion—starting from a kernel at the soma position the algorithm recursively extrudes the neuronal branches. The soma is then scaled to reach the first-order branches.
4. Surface subdivision—the mesh is subdivided to obtain the smooth surface appearance.

While the first two steps use the morphological point representation as data structure, the mesh extrusion builds the control mesh as a polygon-based representation with a list of square-faces referencing four vertices. This coarse model is subdivided and triangulated into a smooth mesh. Finally, in order to be able to relate the new mesh representation to the original morphological point representation, a mapping is computed. Specifically, for each vertex, a section identifier and relative position along that section are computed. We describe some options for computing this mapping in Section 4.2 below.

4.1 Mesh Generation Algorithm

4.1.1 Branch Identification

In the morphological skeleton, sections either branch off from the soma, create a branch point where they join with other sections, or terminate. To build the mesh representation of the branch point connections, we employ a branch construction algorithm to identify which child section would serve as the most natural continuation of the parent section. Branches are identified through a recursive procedure that scans the reconstructed cell data. Starting with the root nodes as the first order of the hierarchy, the procedure is called recursively on all child sections of the current section, until reaching the base case, which is a section without children. At each branch junction of the tree, we apply an algorithm to determine which child

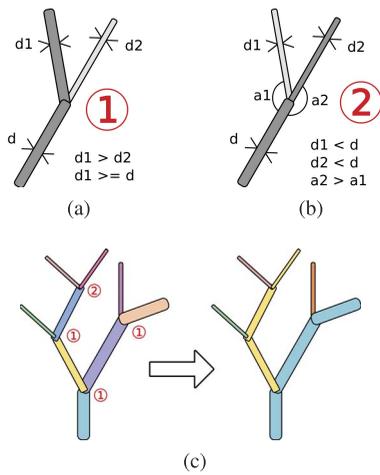


Fig. 4. Identification of the primary branching structures in the neuron morphology. Sections form a branch when either the diameter conditions in (a) or the diameter and angle conditions in (b) are true. (c) Depicts how the original hierarchy based on independent sections is converted into the new hierarchy showing the primary branching structures. Each branch in (c) is marked to indicate the decision made.

section serves as the best continuation of the current branch. The remaining child sections start new branches—although there are typically only two children.

The algorithm prioritizes the section child with the largest diameter. If this diameter is too small compared to the parent section, it chooses the child section with the closest orientation to the parent section (see Fig. 4).

4.1.2 Morphological Point Resampling

The resampling of the original morphological points is essential to generate accurate meshes with a minimum number of vertices. Given that the original data are highly nonuniform, the objective of the resampling is to eliminate redundant vertices while adding vertices at branching points where there are abrupt transitions from the parent to the child sections.

The resampling algorithm handles three conditions (see Fig. 5): branch junctions, redundant points, and sudden diameter changes between the end points of a segment. At branch junctions, three points are necessary to accurately

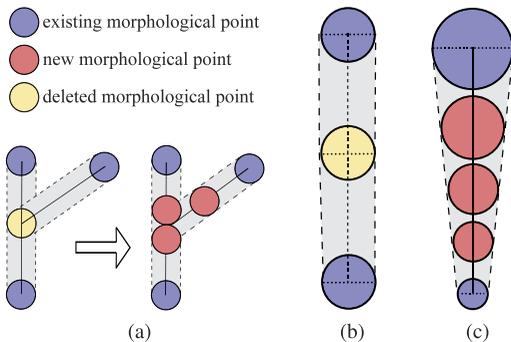


Fig. 5. Sampling of the original morphological points. (a) Adding two points on the parent branch and one on the child branch to facilitate branch connection and constrain the subdivision. (b) Subsampling intermediate morphological point depending on the alignment and the diameter difference. (c) Oversampling between points to constrain the diameter variation between two consecutive points.

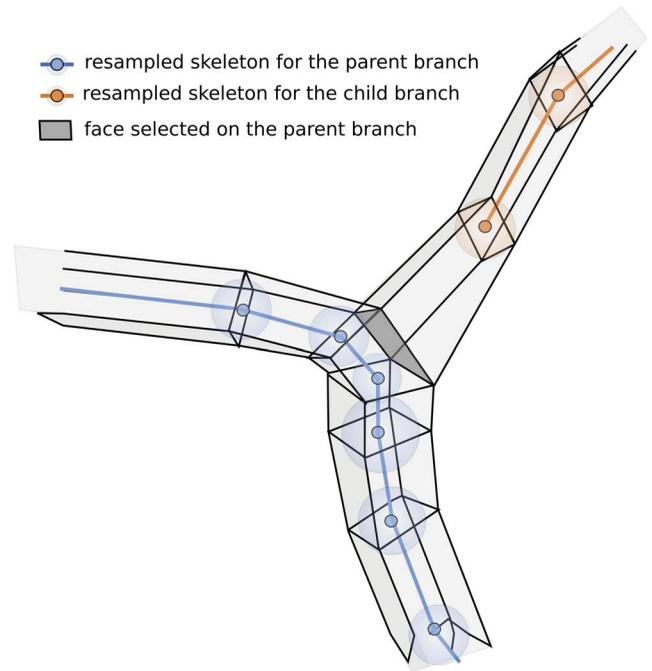


Fig. 6. Extrusion of the square branches along the morphological skeleton. At each morphological point, the extruded square face is scaled to fit the diameter and rotated to become an angle bisector for the angle defined by the previous and the next segment. For each branch junction, the algorithm selects from the parent branch a suitable face oriented toward the child branch and connects it to the new branch.

define the branch connection. First, two points are placed on either side of the original branch point separated by the diameter of the parent branch. The original branch point is then deleted. A third point is added to the child branch at a distance computed (based on branch diameters and angle) to ensure that it is not in the parent branch (see Fig. 5a). To address redundant points, a subsampling algorithm checks the alignment and the diameter variation to delete redundant intermediate points (see Fig. 5b). Upsampling is also necessary when the diameter of one end point from a segment is significantly larger than the other end). In this case, we add additional points with interpolated diameter values to preserve accuracy during subdivision (see Fig. 5c).

The resampling process offers a simple and efficient way to maintain a reasonable trade-off between the accuracy and the final triangle count of the resulting mesh. It has a direct effect on the result of both the branch extrusion and subdivision algorithms.

4.1.3 Control Mesh Extrusion

In this section, we describe the algorithm that extrudes the coarse square mesh from the morphological skeleton (see Fig. 6). This coarse square mesh will be used as a control mesh for the final subdivision surface. The algorithm starts by placing a spherical kernel (polygonal sphere with 36 faces) at the soma. Each first-order branch is then exclusively associated with a kernel face depending on the face orientation and the position of the first point of the branch. This state constitutes the base for the recursive branch extrusion process. A full branch with all the sections it comprises is first processed before proceeding to its children branches.

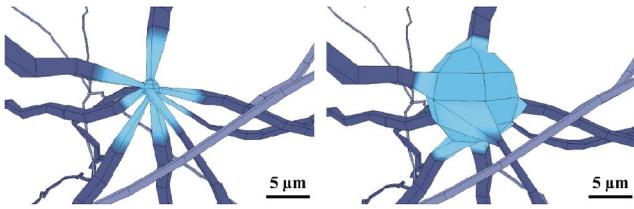


Fig. 7. Soma construction. This picture introduces two fundamental steps in the construction of the soma. In the first part (a), the kernel is placed and connected to the branches. In the second part (b), the kernel is scaled to reach the closest branch.

To build the control mesh for the subdivision surface, we used a square-faced prismatoid extrusion: the copy of an initial convex quadrilateral face is translated following the branch path to the next point and then connected to the previous face to form a square-faced prismatoid. The algorithm performs successive extrusions until the morphological point sequence of the branch is completed. The diameter of the morphological points is preserved by performing orientation adjustments and rescaling the extruded face. Once a branch is extruded, the same operation is recursively applied to the child branches (see Fig. 6).

During the extrusion process, the choice of the initial face for child branches determines the level of difficulty in preserving the correct diameter of the proximal part of the child branch at the connection site. On one hand, starting from a unit square at the location of the first child point next to the branch point simplifies the diameter fitting, however, this implies a proper alignment between the vertices of a child branch and those from the parent branch.

On the other hand, starting directly from a quadrilateral face on the parent branch facilitates the connection between the parent and child branches, but induces a series of geometrical operations to convert the extruded faces into the scaled squares that fit the diameters of the morphological points. In the current implementation of the algorithm, the branch is extruded from a unit square located at the position of the first morphological point of the child after the branch point, and the vertices of the first child branch face are connected to a side of the parent branch by minimizing the torsion of the quadrilateral prism connecting both faces.

Once the control mesh is extruded, the kernel is scaled to match the actual soma size (see Fig. 7b). To determine the correct size of the soma, we compute the distance between the center of the kernel and the closest first-order branch. This distance is scaled by a weighting coefficient (<1) to ensure that all branches are outside the final soma shape.

4.1.4 Surface Subdivision

After combining all branches into a single structure and connecting the first branches to the soma, we arrive at a complete, yet coarse, polygonal representation of the neuron with the correct connections between the branches and the soma. The polygonal surface is closed due to the successive extrusion maintaining the continuity of the mesh. This mesh is then used as a control mesh to generate a subdivided surface. To smooth the mesh surface, we use the first iteration of Catmull-Clark subdivision [17], which preserves the

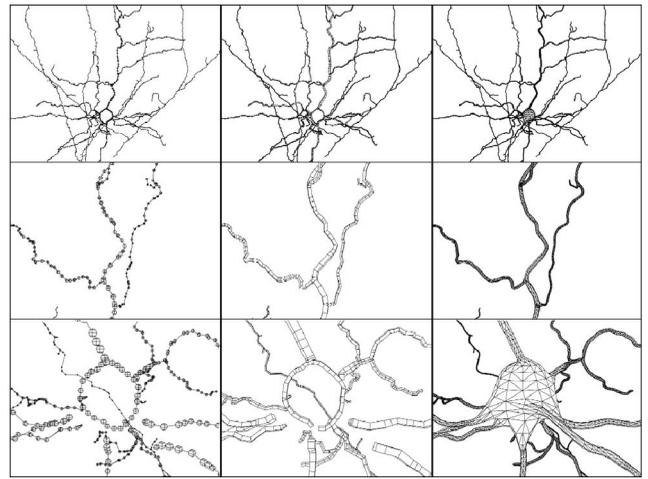


Fig. 8. Starting from the original skeleton of the morphology with morphological points (a), the algorithm extrudes a square simple shape (b). The neuronal shape is smoothed using the first iteration of the Catmull Clark algorithm (c).

diameters of the branches (see Fig. 8c). Other nonshrinking subdivision techniques could also be used here [34].

Once the smoothing algorithm is completed, a triangulation transform is applied to the quad polygons to facilitate the texture mapping and the computation of the surface normals. Triangle strips can then be computed from the final mesh to reduce memory consumption and potentially improve vertex cache efficiency of the final mesh. Depending on the implementation of the smoothing algorithm, the rendering performance can be greatly improve due to the lighting calculations and vertex transformations saved [10].

4.2 Mapping Membrane Mesh Vertices Back to the Morphological Skeleton

The resulting mesh from the process described above is already sufficient for a wide range of applications, including structural visualization and evaluation of spatial relationships between neurons. However, being able to relate mesh vertices to the original morphological skeleton is a requirement for other visualization and simulation applications. We therefore propose a mapping method that associates with each vertex, a section identifier (from the original morphological representation) and a normalized relative distance into the section. This mapping can be created either during mesh generation or later, as a postprocessing step. Fig. 9 shows how these mapping attributes can be used to map values from a morphologically based model to the mesh vertices.

To obtain the mapping during the mesh generation it must be noted that a vertex belongs uniquely to one section, which is an important consideration for the subdivision step. Such a process would start assigning initial values to the morphological points. During the resampling step the attributes are updated accordingly and at path extrusion, newly created vertices inherit the attributes from the originating vertices. Finally, in the surface subdivision step the attributes need to be interpolated/inferred. As already pointed out, section identifiers cannot be interpolated; therefore standard texture coordinate interpolation does not apply. One possible approach would be to first assign

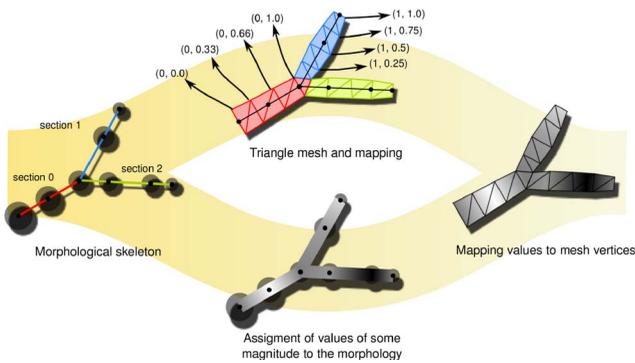


Fig. 9. Mesh to morphological skeleton mapping. Some types of structural or functional neuron models are obtained from the morphological skeleton. The parameters or simulation results of these models can be applied as rendering attributes such as color onto the mesh. In order to do so, a mapping from the mesh vertices to the skeleton is required, which relates a vertex with a section and normalized relative position within that section.

attributes to faces by vertex aggregation. In the Catmull-Clark subdivision, any final quadrilateral can be backtracked to a single parent base face. Thus, the section mapping can be inherited and the relative position interpolated during subdivision iterations. The final vertices can then be assigned attributes interpolated from the faces touching them, assigning section identifiers according to user selected criteria.

Another method for computing the mapping is to do it independently from the mesh generation as a separate postprocessing. The decoupling of both calculations provides flexibility to use existing tools in the implementation of the mesh generation and also gives freedom for changes. Since the efficient algorithm based on nearest neighbor queries using an octree avoids all the issues involving discrete section identifiers, we choose this approach to relate mesh vertices to the original morphological skeleton.

Once all the information (vertices locations, triangles indices, triangle strips, and mapping information) are computed, the data are combined together into a single binary file, corresponding to the original morphological skeleton.

5 RESULTS

Using the methods described in this paper, we have generated 173 unique neuron membrane meshes. Given the variety of neuron shapes, we have seen that the mesh generation algorithm generalizes well to a broad range of neurons as found in the neocortex (see Fig. 10). Our algorithm is not only generic enough to be used in a batch process to correctly generate meshes without adjusting the parameters, but also robust enough to support inconsistencies in the data like shifted branches or multiple child sections (up to five child sections at a branch point). However, the algorithm requires that the sections are organized hierarchically and does not support cyclic graphs.

5.1 Validation

To validate the generated meshes, we compare the mesh branches with the original biological measurements. This validation exploits the mapping information to compare the

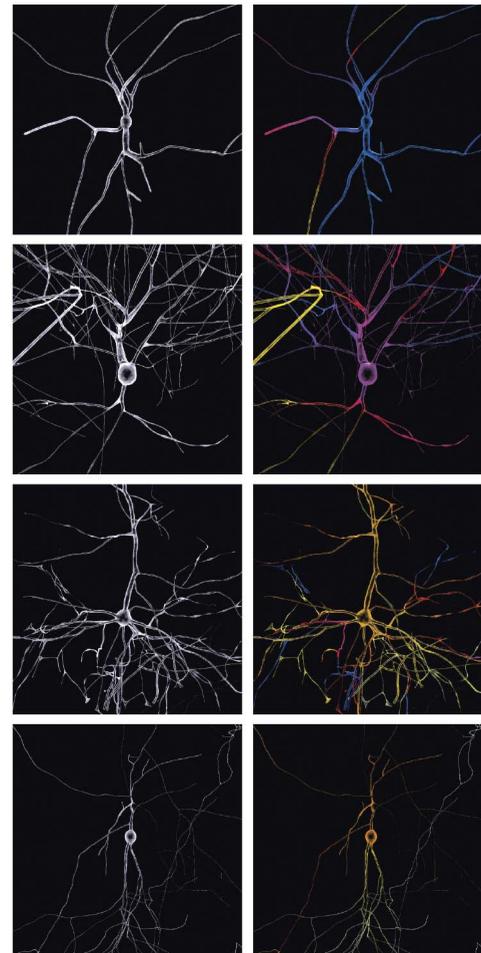


Fig. 10. Mesh generation for different morphologies. This mesh generation algorithm can be applied to a wide variety of cortical neuron types. The mapping information joined with the mesh representation allows the coloring of vertices depending on electrophysiological simulations.

position and the diameter of the mesh branches with those in the original morphology.

Using mesh point clusters—circular groups of vertices derived from the subdivision process—we compute points and diameters that correspond to the morphological points of the original biologically measured skeleton. The center of mass and the extent of each mesh point cluster provide the position and diameter information for any given mesh location, which can then be compared with the original morphological data (see Fig. 11). Using the mapping and linear interpolation between the morphological points, we find an interpolated morphological point for each mesh point cluster. The mesh point cluster is defined as a set of mesh points located on a plane perpendicular to the segment (see Fig. 11b). The difference between the diameter of the mesh and the diameter of the original point evaluates the shrinking effect due to the subdivision algorithm (see Fig. 11a). The results are clustered around zero with a maximum absolute difference equal to 0.5 percent (see Fig. 11c). The alignment of the mesh model with morphological skeleton is assessed by computing the out-of-axis distance of the mesh clusters projected in a common referential for all the morphological points (see Fig. 11b). Each center of mass of the mesh point cluster (for example, $c1$ and $c2$) is projected on a plane,

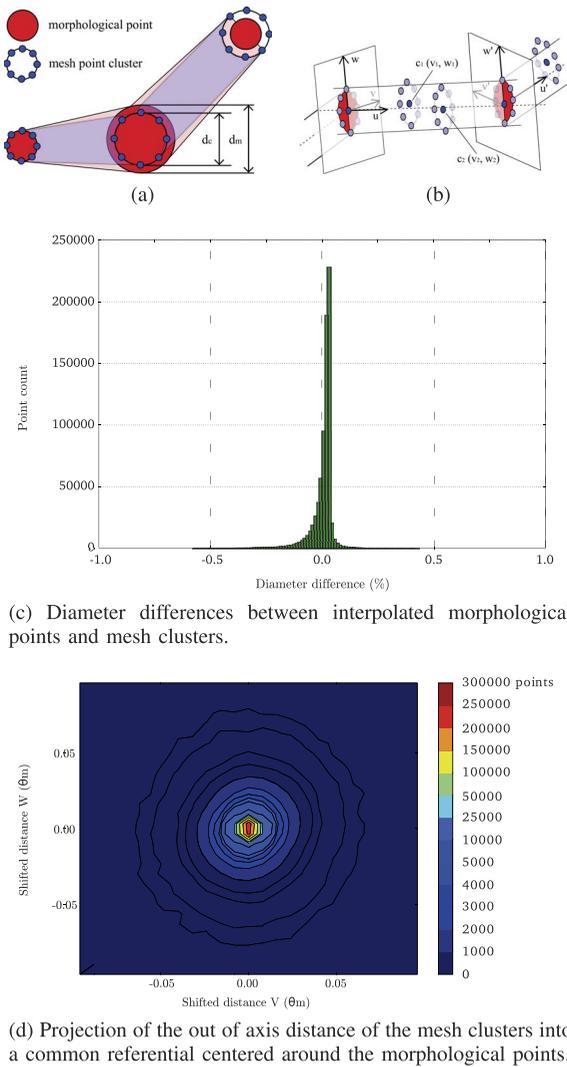


Fig. 11. Mesh accuracy algorithm and results for 173 neurons with a combined total of 802,357 mesh point clusters. The diameter variation distribution (a) demonstrates that there is no shrinking effect due to the subdivision (c), while the symmetric distribution of out-of-axis distance points (b) reveals that nearly all absolute distances between the mesh point clusters and the morphological points are less than 0.03 micron (d).

perpendicular to the segment direction. The consistency between the projection planes is maintained by a recursive rotation starting from the global referential. Thus, if q is the quaternion to express the rotation from u to u' , we have $v' = q \dot{u} v$ and $w' = q \dot{u} w$. The result shows a symmetric distribution of the cluster projection around zero (see Fig. 11d). More than 98 percent of the points are located at less than 0.03 micron from the corresponding morphological point. This error corresponds to one tenth of the actual microscopic resolution during the reconstruction process. Thus, the error introduced during the mesh reconstruction is negligible compared to the precision of the skeleton extraction. The specific form the mesh takes at the branch junctions cannot be directly validated because the morphological information is too coarsely sampled to provide the detailed curvature information. However, by validating the morphology points that define the incoming and outgoing sections, we ensure that the angles of the branch junction are preserved in the mesh.

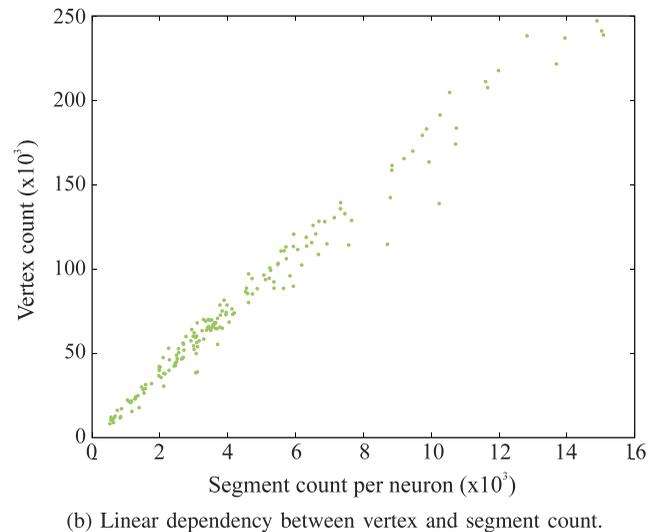
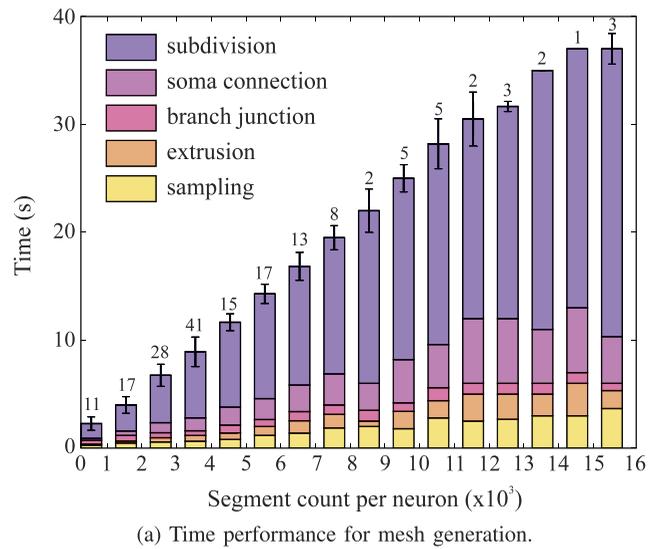


Fig. 12. Performance of the mesh generation computed on a standard 2 GHz Core Duo PC with 2 GB RAM using a set of 173 neurons with different morphological complexity. (a) shows a linear dependency between the morphological complexity (segment count) and the mesh generation runtime. (b) illustrates the linear correlation between the segment count and the amount of vertices generated.

5.2 Measuring Performance

To illustrate the mesh generation performance, we benchmarked the algorithms using a set of neurons with a range of different sizes and complexities (see Fig. 12). The runtime scales linearly with the section count as expected. The most expensive part of the mesh generation algorithm is the smoothing of the coarse mesh using the subdivision surface algorithm (see Fig. 12a). Fig. 12b demonstrates that the number of vertices generated scales linearly with the complexity of the neuron measured by the segment count.

6 VISUALIZATION OF ELECTROPHYSIOLOGICAL SIMULATIONS

Visualizing single neurons and networks of cells require a flexible neuron representation that provides accurate depictions of neuron morphologies with minimal artifacts. As both single cells and networks can be viewed at different

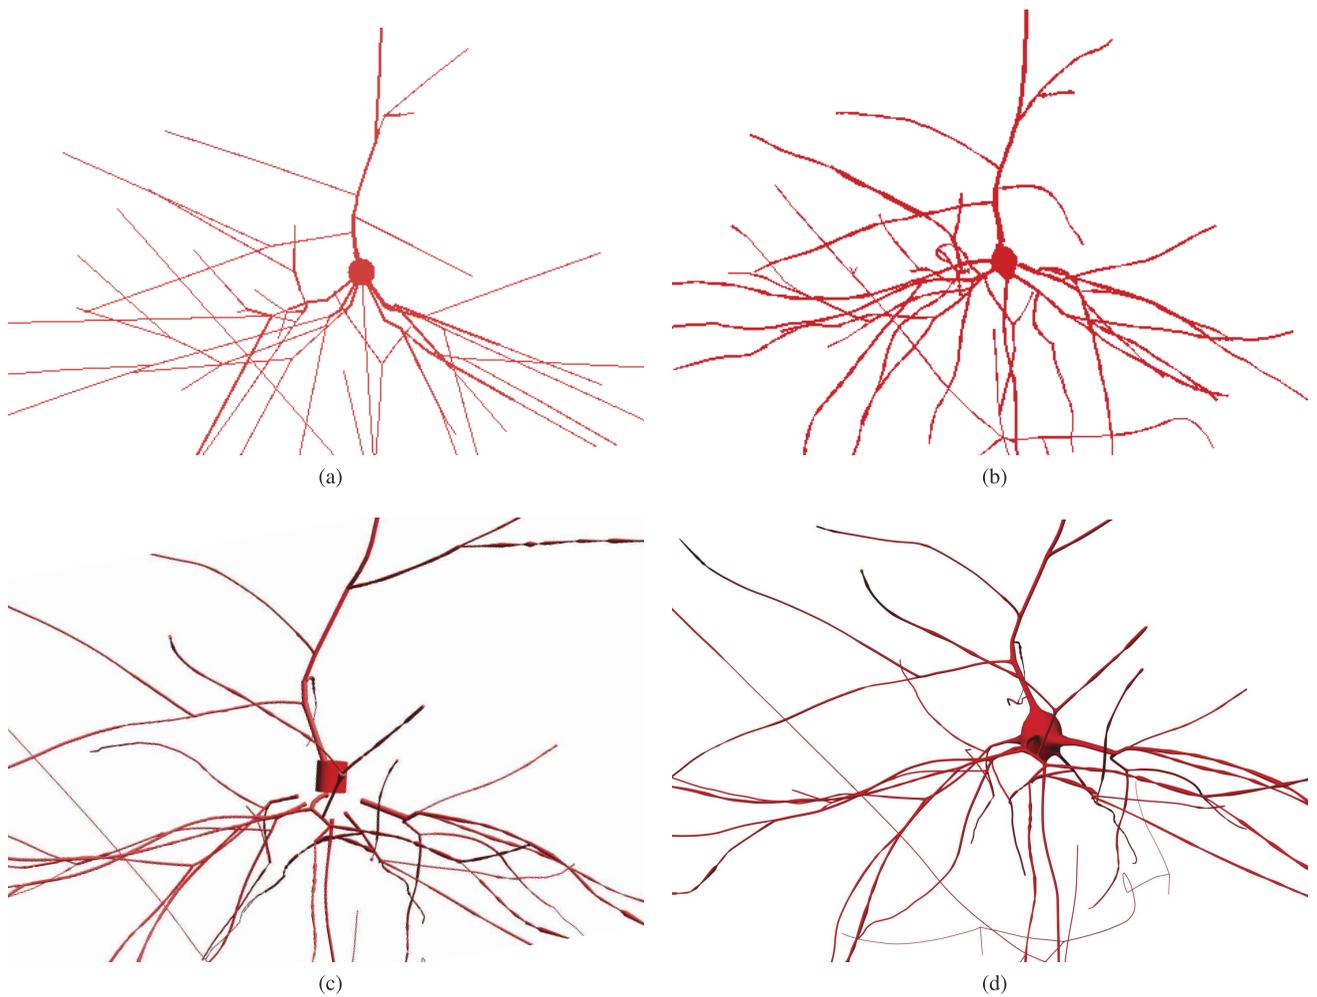


Fig. 13. Comparison of the state-of-the-art neuron representations with the new mesh representation. (a) the line-based representation used in Genesis where each line represents an electrical compartment of the simulated neuron. (b) a line-based representation in NEURON where each line corresponds to a morphological segment. (c) the tube-based representation from Neuroconstruct converts each morphological segment into a continuous faceted pipe. (d) the new continuous mesh representation for the same original morphological skeleton.

viewing distances, for example, to examine the activity of a single neuron in the context of large-scale network activity, the representation should permit artifact-free rendering at multiple scales. In addition, this representation should be suitable for high-quality renderings for scientific publications and communication to a larger public audience. Here, we compare the new mesh representation with existing visualization techniques. First, we compare the current state-of-the-art techniques from well-known simulator software packages. We then examine specific issues in visualizing single cells and point out rendering artifacts that are reduced by using a mesh representation. Finally, we demonstrate that for visualizing electrical simulations of large networks of neurons, meshes provide increased flexibility of visualization, enabling new questions to be answered in the study of complex network dynamics.

6.1 Comparison with State of the Art

The current state of the art for visualizing neuron models and simulation results is represented by three software packages: Genesis [39], NEURON [5], and Neuroconstruct [14], and from neuron reconstruction packages like NeuroLucida [13]. We therefore compared the graphic

representation of neurons provided by each of the above packages (see Fig. 13). Genesis provides a graphical representation of only the electrical cable-equivalent compartments for a given morphology (which are a subsampled portion of the actual morphological skeleton). NEURON produces a line plot of a given neuron using the actual morphological points including diameter variations, thus providing a complete two dimensional representation of the neuron. The soma is represented as a flat disc of the measured somatic diameter. Neuroconstruct produces a three-dimensional tube-based representation of a neuron using morphological point data and the variation of diameters. The soma is represented by a cylinder with the connected branches displayed adjacently. The new technique, described in this paper, produces a continuous mesh representation that fits the morphological points and diameter variations. The mesh connects all branches to the soma through interpolation while maintaining the original measured soma diameter.

6.2 Visualization of a Single Cell Electrophysiological Simulation

The structure and form of a neuron is intrinsically linked to its function and interactions with other neurons. The spatial

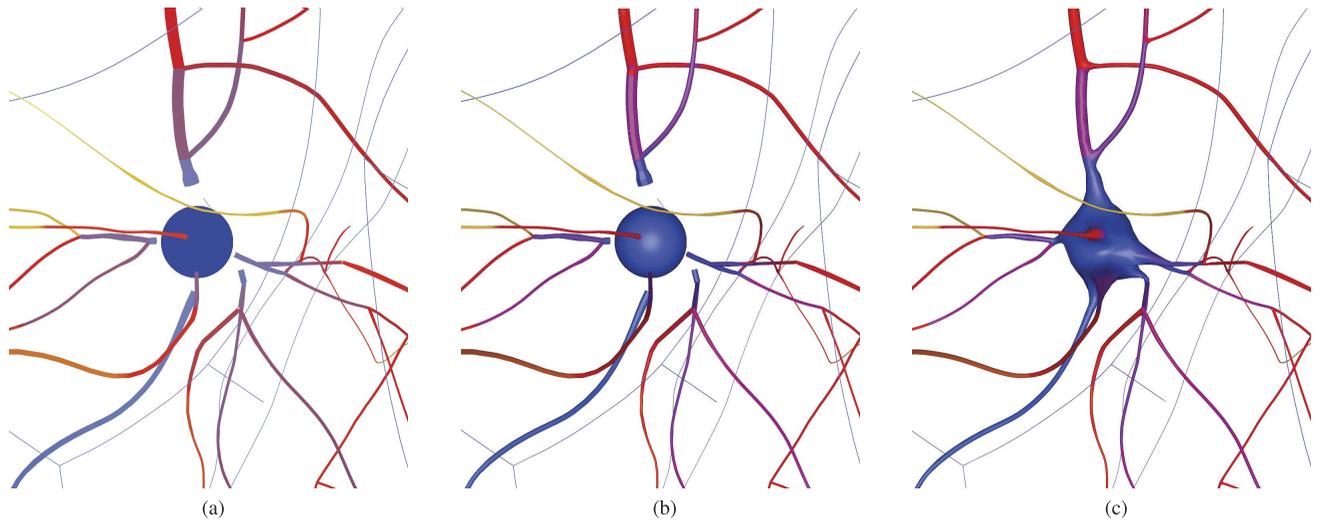


Fig. 14. Three representations for a single simulated neuron. The line-based representation depicted in (a) misses the shading and the occlusion clues that are crucial for the perception of complex arborized structures. Both tube-based and mesh representations, respectively, in (b) and (c), exposes shading and occlusion clues. However, in the case of the mesh representation the connection between the first-order branch with the generated soma raises the contour clue and facilitates the overall perception of the cell.

structure of the dendrites forms the typical region for incoming electrical input. The spatial location and propagation of electrical potentials on the neuronal morphology is key to understanding how a single neuron integrates signals from other neurons and ultimately makes the decision to “spike” and signal other neurons [18]. It is therefore essential to have clear cues about the three-dimensional form and relative distances of the different branches and bifurcations of dendritic morphology from the soma when visualizing single neuron electrical activity. Depth cues aid the perception and description of a three-dimensional scene in terms of distance, surface orientation, curvature, shape, or form. The advantages of using the depth cues are as follows, stereo perception improves the depth map, shading improves the curvature and surface orientation, and occlusion improves the depth ordering [4]. It is well understood that shading can improve the perception of solid shapes [24], [35] and in particular can be helpful to visualize the arborized complex structure of neurons. In the case of a line rendering of a single neuron, (see Fig. 14a) shading or reliable occlusion information is unavailable, making it difficult to discern distal or proximal portions of a dendrite or axon relative to the soma. Tube-based rendering (see Fig. 14b) allows shading to be employed to provide depth cues and improve the perception of the three-dimensional structure of the neuron and the relationship of axons and dendrites to the soma. The new mesh representation provides both depth and occlusion cues (see Fig. 14c) similar to the tube-based representation, but the mesh forms a continuous connection between the first-order branches and the cell soma.

6.2.1 Artifacts from Translucency

Translucency can play an important role in providing an overview and a detailed visualization of phenomena of interest. View-dependent transparency has been long used in technical illustration to reduce clutter and highlight specific features or a model, and more recently has been employed for computer-based illustration [8], [37], [9].

Here, we apply a simple rule that combines structural information (the diameter of the axon or dendrite fiber) with simulated voltage data to define the translucency of the fibers.

Both the mesh and tube-based representations were used to explore translucency combined with depth cues. However, due to the fact that the tube-based representation is composed of joined cylindrical elements, it exhibits artifacts at each branch point when translucency is employed (see Figs. 15a and 15b). On the other hand, the mesh representation is continuous and forms artifact free transitions at branch points and at junctions with the soma (see Figs. 15c and 15d). Therefore, in order to employ translucency to enhance visualization, a continuous artifact-free representation such as the new mesh representation, is required.

6.3 Visualization of a Simulated Network

Visualization of a large population of neurons represents a particular challenge given the density of complex axonal and dendritic fibers that pervade the scene. One strategy to reduce the overall density and complexity of the scene is to use translucency to highlight the activity of interest in the dense neuronal network. The translucency rule (described in the previous section) serves to highlight cell activity when the voltage deviates from a baseline level (for example, when synaptic input from another cell arrives on the dendrite or the cell spikes and communicates with other neurons). By eliminating the density of the neuronal fibers that are not actively participating in the network at a specific time, visual clutter is greatly reduced and patterns in the complex network activity are more readily discerned as shown in Fig. 16. Specifically, we examine all the inhibitory presynaptic neurons for an individual layer V pyramidal cell in a model of rat neocortex. Inhibitory neurons typically inhibit spike generation in a layer V pyramidal cell. When any of the inhibitory cells are active they will prevent the excitatory layer V pyramidal cell from firing through their influence on its electrical activity. Different groups of inhibitory cells can be active at

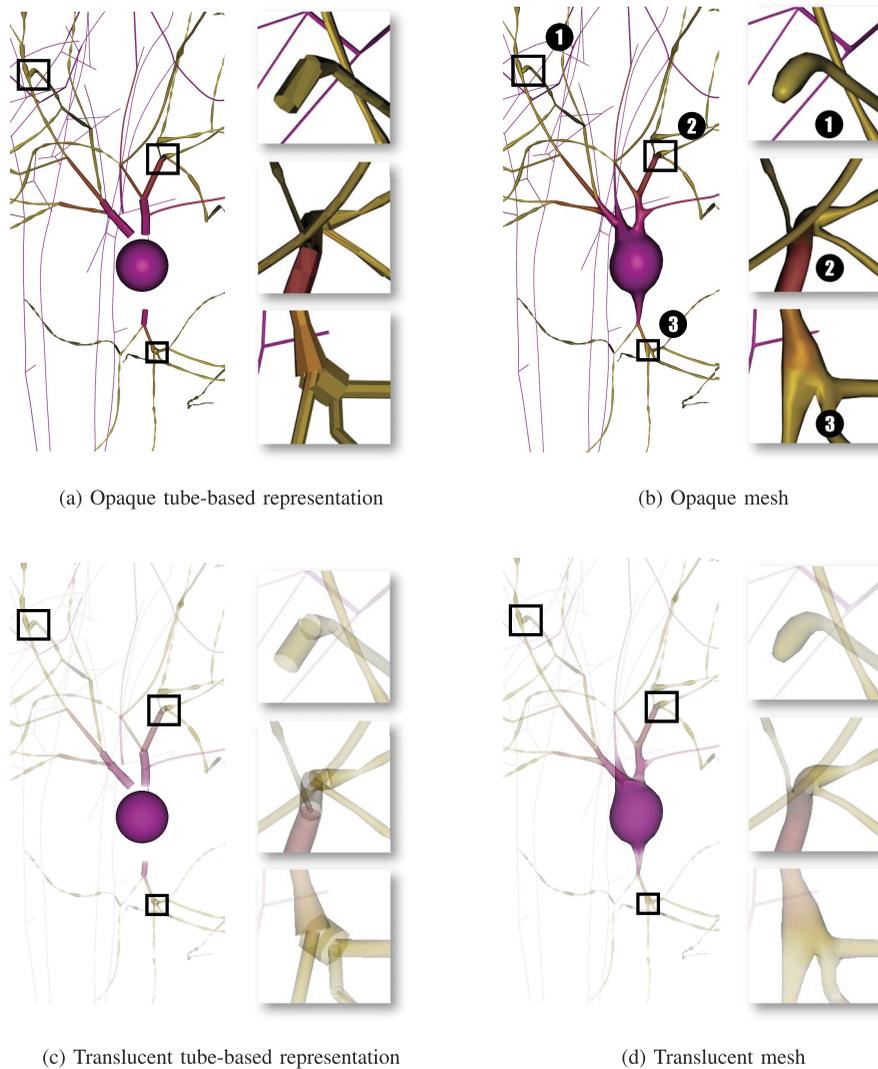


Fig. 15. Artifacts in the mesh and tube-based representations in opaque and translucent rendering modes. Three locations are observed on the different representations. The first location (1) exposes the termination of a dendrite with a strong curvature and increased diameter, whereas the two other locations (2 and 3) show complex branch points. The opaque tube-based representation (a) and the opaque mesh (b) compose the first row. In this case the coarse tube-based representation competes with the mesh representation despite the continuity issue between the first-order branches and the soma. The second row exposes the same representations rendered with translucency. While the continuous mesh membrane (d) does support translucency without severe artifacts, the tube-based representation (c) exposes the limits of its tube-based representation and makes difficult the interpretation of the simulation color.

different times—often they dynamically shift and oscillate. Therefore, translucency provides a technique to show only the population of inhibitory cells active at a given time. In addition, using this technique enabled us to see the detailed correlational structure of voltages in independent somas and dendritic fibers during network activity. Remarkably, we saw that the voltages in the dendrites of independent neurons could be highly correlated in space during different phases of network activity. This provides insight into the dynamic organization of excitatory and inhibitory networks in neocortex, which is a central question in neuroscience [15], [27].

7 DISCUSSION

In this paper, we describe a process and series of algorithms to transform morphological point reconstruction data from

experimentally traced neurons into a continuous membrane that is true to the original neuron while preserving the mapping between the original sampled neuron and the polygons of the mesh.

7.1 Versatile Neuron Membrane Mesh Representation

The process described in this paper contributes a novel representation for simulation-based research in neuroscience. This new mesh representation is virtually artifact free, which facilitates advanced rendering techniques and scientific exploration without distracting artifacts or clutter. We have demonstrated that this membrane representation provides important 3D depth cues from shading, occlusion, and contour continuity. In addition, we used these meshes with stereoscopic rendering to provide immersive visualization experiences. The meshes also currently serves as the

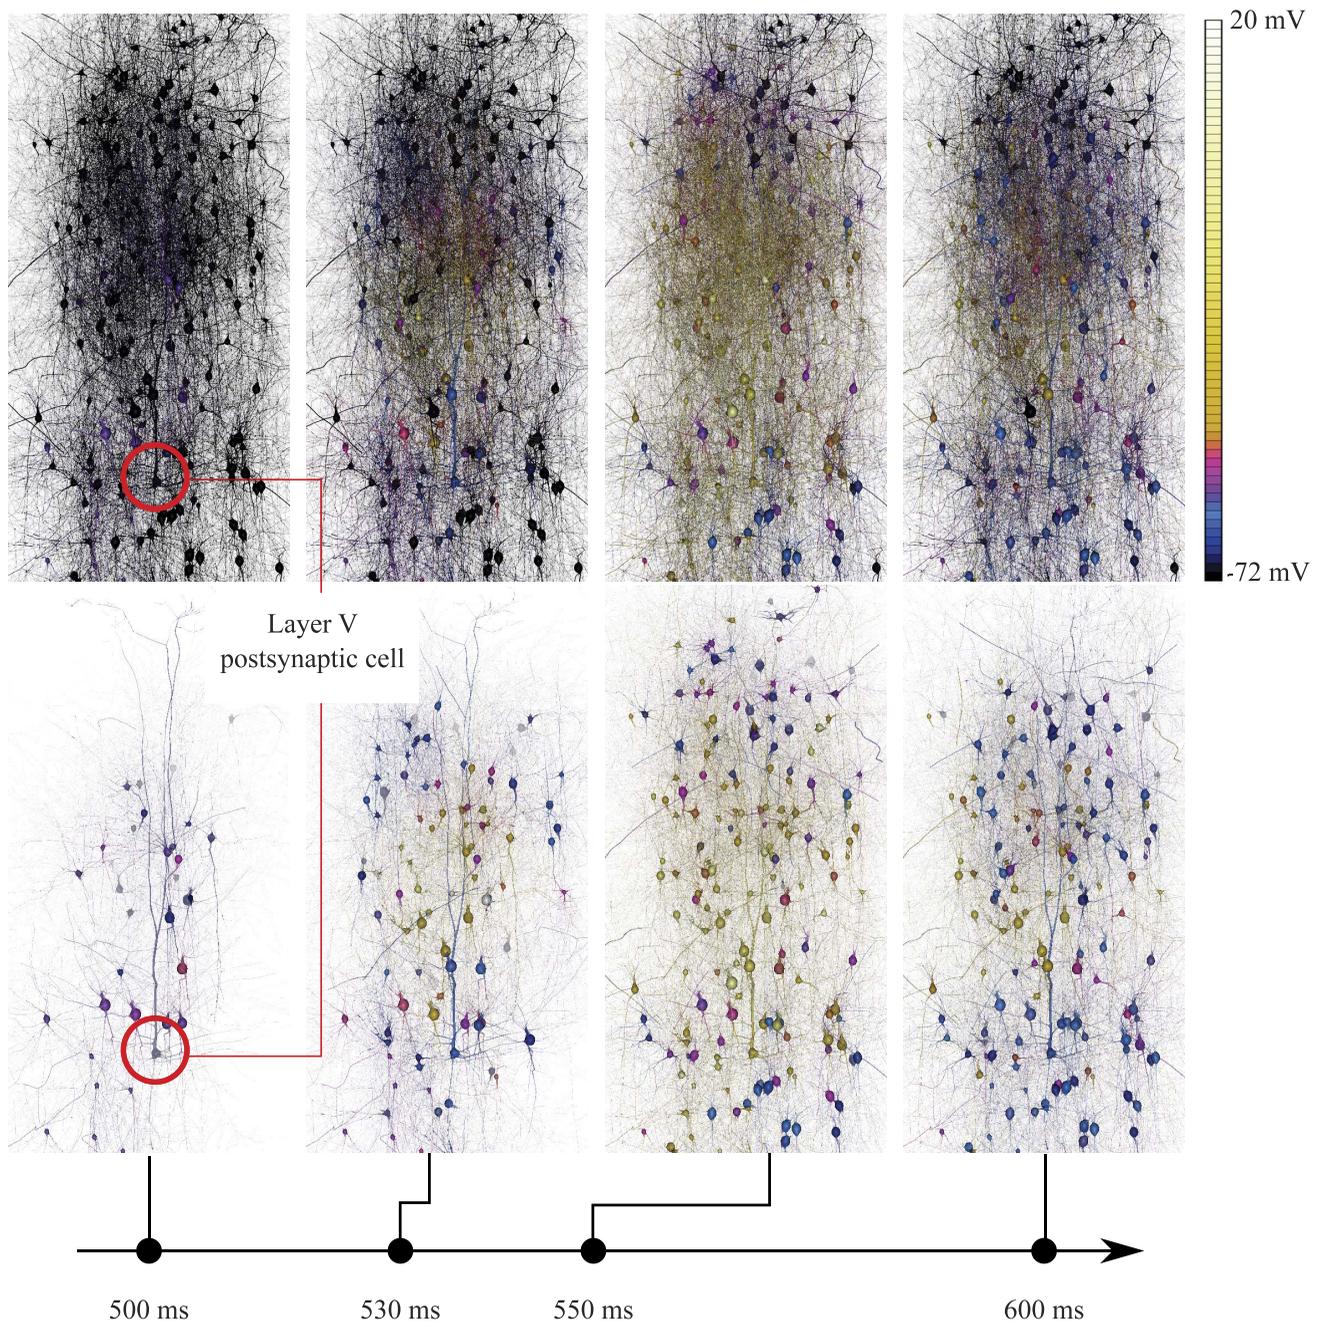


Fig. 16. Visualization of dynamics of a single layer V pyramidal cell and its presynaptic inhibitory cells. A complex neuronal microcircuit is simulated using large-scale compartment-based electrical simulation and a single specific postsynaptic neuron is extracted with all its presynaptic inhibitory cells. The figure depicts four time stamps of the simulation, where the overall network is first at a resting potential, then gets activated, and finally reverts to a stable state. Membrane voltage around -72 mV is considered as hyperpolarized (or inhibited). Around -40 mV, the membrane is depolarized (or excited). After a certain threshold of depolarization at the soma, the neuron will generate a “spike” and convey an electrical signal to its postsynaptic cells. In the first row, the behavior of the studied population is rendered without translucency: all the cells are visible and it is difficult to distinguish which presynaptic inhibitory cells are involved in the activity of the postsynaptic cell. In the second row, a translucency factor is applied to the cells: the opacity of the branches is defined by the combination of the thickness of the branch and the simulated membrane potential. In this case, only the inhibitory cells actively involved in the inhibition of the pyramidal cell are apparent. The correlation between the opacity and the thickness of the branches reveals the primary branches of the cells and removes the smaller branches, which would appear as aliasing at this distance and would add only noise in the rendering. Note that the dendritic and somatic voltages of independent neurons can be highly correlated in both space and time.

structural basis for more detailed models that take into account molecular reactions and diffusion throughout the internal structure of the neuron [31]. They are also used for the visualization of a variety of neuronal properties including ion channel densities and currents, and synaptic properties.

7.2 Limitations

The accuracy of the mesh was validated by comparing the diameter and the position of the vertices in the branches against the diameter and position of the morphological point representation. The robustness of this technique has been demonstrated by generating nearly 200 neuron meshes

from a broad variety of available reconstructed neurons [16] without requiring the adjustment of any fitting parameters.

However, the original morphological data suffer from low resolution measurements of the soma and branch junctions. Although bright field microscopy reconstructions remain relatively inexpensive and prevalent, laser-scanning confocal or electron microscopy (EM) could ultimately offer high resolution data to accurately capture the soma and branch junctions in the future.

7.3 Lessons Learned

In the course of this work, we have seen that geometrically and topologically accurate depictions of the neuronal morphology are much more than simple eye candy. The spatial structure is extremely relevant to the dynamics and function of single neurons. Accurate visualization of the spatial and temporal aspects of electrical activity within individual neuron morphologies is therefore essential to the scientific discovery process. Neurons themselves are difficult entities to visualize. The ratio between the finest fibers ($<1 \mu\text{m}$) and the overall size of a neuron (up to 1 mm) makes the rendering process challenging. The overall quantity of data to accurately represent all portions of the neuron is substantial and with large networks composed of unique neuron meshes, the IO performance is a significant challenge that remains to be addressed.

Another challenge that had to be addressed for high resolution renderings was to combine over- and multi-sampling techniques in order to accurately antialias the images. Multisampled meshes is a straight forward and ubiquitous hardware feature, however, in our experience, some tube-based rendering techniques are not amenable to multisampling. The reason is that, except for recent hardware that allows a full shader execution per sample, discard fragment operations kill all fragment samples, thus quadratic equation-based or ray-casted primitives from the literature are multisampling incapable. On top of the technical issues, we had to learn how to use such complex neuronal morphologies to visualize network simulations. Here, it became clear that translucency was essential to avoid single cell artifacts and reduce scene clutter. The density of neuronal fibers is such that most activity will be obscured if some translucency is not employed to reduce the density of the scene.

We saw that the tube-based representation can be sufficient for many purposes. If the scene does not involve close views of the network, tube-based representations are virtually indistinguishable from those generated from a mesh. However, the mesh representation becomes essential when viewing single neurons or networks at multiple spatial scales. This mesh provides a high quality artifact free representation across all scales. We have found that it is possible to perform near real-time rendering of simulation data on 100 meshes of around $\sim 150\text{k}$ triangles on modern PC hardware. Additional optimizations, such as levels of detail and the appropriate culling algorithm, may enable even larger scenes to be interactively explored.

7.4 Insights

An important scientific insight was gained through these visualizations—that somatic and dendritic voltages of

independent neurons could be highly correlated during network activity. The mechanisms producing such correlational structure are under active investigation. In general, high fidelity rendering with a minimum of artifacts allows scientists to visualize detailed physiological simulations and deepen the understanding of the complex neuronal dynamics of both individual neurons and large-scale networks.

7.5 Future Directions

We have shown that our approach generalizes quite well for different classes of cortical neurons. The meshes could also serve as the structural basis of more detailed neuron models that take into account molecular reactions and diffusion throughout the internal structure of the cell [31]. We intend to apply and generalize the algorithm to handle other neuron classes from throughout the brain in the near future. In addition, this method could also apply to glial cells—important support cells for neurons—as well. The vasculature in the brain—the complex, arborized structures of the blood supply—could also be modeled using this technique.

ACKNOWLEDGMENTS

This work was supported by the Blue Brain Project and the Cajal Blue Brain. We thank Alexandre Chavaillaz for his contributions during the initial stages of this work, Thomas Berger for the single cell reconstructions used in this study, Rajnish Ranjan for simulation support, and Srikanth Ramaswamy for reviewing the manuscript. The authors are grateful to the reviewers for their constructive comments.

REFERENCES

- [1] P. Alliez, G. Ucelli, C. Gotsman, and M. Attene, "Recent Advances in Remeshing of Surfaces," Part of the State-of-the-Art Report, AIM@SHAPE EU Network of Excellence, 2005.
- [2] U.S. Bhalla and J.M. Bower, "Exploring Parameter Space in Detailed Single Neuron Models: Simulations of the Mitral and Granule Cells of the Olfactory Bulb," *J. Neurophysiology*, vol. 69, no. 6, pp. 1948-1965, June 1993.
- [3] K. Briggman and W. Denk, "Towards Neural Circuit Reconstruction with Volume Electron Microscopy Techniques," *Current Opinion in Neurobiology*, vol. 16, no. 5, pp. 562-570, Oct. 2006.
- [4] H.H. Bulthoff and H.A. Mallot, "Integration of Stereo, Shading and Texture," *AI and the Eye*, pp. 119-146, 1990.
- [5] N.T. Carnevale and M.L. Hines, *The NEURON Book*. Cambridge Univ. Press, Feb. 2006.
- [6] M. Dailey, G. Marrs, J. Satz, and M. Waite, "Concepts in Imaging and Microscopy Exploring Biological Structure and Function with Confocal Microscopy," *Biological Bull.*, vol. 197, no. 2, pp. 115-122, Oct. 1999.
- [7] W. Denk and H. Horstmann, "Serial Block-Face Scanning Electron Microscopy to Reconstruct Three-Dimensional Tissue Nanostructure," *PLoS Biology*, vol. 2, no. 11, pp. 1900-1919, Nov. 2004.
- [8] J. Diepstraten, D. Weiskopf, and T. Ertl, "Transparency in Interactive Technical Illustrations," *COMPUTER GRAPHICS FORUM*, vol. 21, pp. 317-326, 2002.
- [9] N. Elmqvist, U. Assarsson, and P. Tsigas, "Employing Dynamic Transparency for 3D Occlusion Management: Design Issues and Evaluation," *Human-Computer Interaction - INTERACT '07*, In Cécilia Baranauskas, Philippe Palanque, Julio Abascal, and Simone Barbosa, eds., pp. 532-545, Springer, 2010.
- [10] F. Evans, S. Skiena, and A. Varshney, "Optimizing Triangle Strips for Fast Rendering," *Proc. Seventh Conf. Visualization '96*, pp. 319-326, 1996.
- [11] J.C. Fiala, "Three-Dimensional Structure of Synapses in the Brain and on the Web," *Proc. Int'l Joint Conf. Neural Networks (IJCNN '02)*, vol. 1, pp. 1-4, 2002.

- [12] S.F. Frisken, R.N. Perry, A.P. Rockwood, and T.R. Jones, "Adaptively Sampled Distance Fields: A General Representation of Shape for Computer Graphics," *Proc. 27th Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 249-254, 2000.
- [13] J.R. Glaser and E.M. Glaser, "Neuron Imaging with NeuroLucida-a PC-Based System for Image Combining Microscopy," *Computerized Medical Imaging and Graphics: The Official J. Computerized Medical Imaging Soc.*, vol. 14, no. 5, pp. 307-317, Oct. 1990.
- [14] P. Gleeson, V. Steuber, and R.A. Silver, "Neuroconstruct: A Tool for Modeling Networks of Neurons in 3D Space," *Neuron*, vol. 54, no. 2, pp. 219-35, Apr. 2007.
- [15] B. Haider and D.A. McCormick, "Rapid Neocortical Dynamics: Cellular and Network Mechanisms," *Neuron*, vol. 62, no. 2, pp. 171-189, Apr. 2009.
- [16] M. Halavi, S. Polavaram, D. Donohue, G. Hamilton, J. Hoyt, K. Smith, and G. Ascoli, "NeuroMorpho.Org Implementation of Digital Neuroscience: Dense Coverage and Integration with the NIF," *Neuroinformatics*, vol. 6, no. 3, pp. 241-252, 2008.
- [17] M. Halstead, M. Kass, and T. DeRose, "Efficient, Fair Interpolation Using Catmull-Clark Surfaces," *Proc. 20th Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 35-44, 1993.
- [18] E. Kandel, J. Schwartz, and T. Jessell, *Principles of Neural Science*, fourth ed. McGraw-Hill Medical, Jan. 2000.
- [19] J. Klein, F. Ritter, H.K. Hahn, J. Rexilius, and H.-O. Peitgen, "Brain Structure Visualization Using Spectral Fiber Clustering," *Proc. ACM SIGGRAPH '06 Research Posters*, p. 168, 2006.
- [20] W.E. Lorensen and H.E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," *ACM SIGGRAPH Computer Graphics*, vol. 21, no. 4, pp. 163-169, 1987.
- [21] H. Markram, "The Blue Brain Project," *Nature Rev. Neuroscience*, vol. 7, no. 2, pp. 153-160, Feb. 2006.
- [22] Z. Melek, D. Mayerich, C. Yuksel, and J. Keyser, "Visualization of Fibrous and Thread-Like Data," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1165-1172, Sept./Oct. 2006.
- [23] M. Meyer, R.M. Kirby, and R. Whitaker, "Topology, Accuracy, and Quality of Isosurface Meshes Using Dynamic Particles," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1704-1711, Nov./Dec. 2007.
- [24] E. Mingolla and J.T. Todd, "Perception of Solid Shape from Shading," *Biological Cybernetics*, vol. 53, no. 3, pp. 137-151, 1986.
- [25] V. Petrovic and J. Fallon, "Visualizing Whole-Brain DTI Tractography with GPU-Based Tuboids and LoD Management," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1488-1495, Nov./Dec. 2007.
- [26] H. Qu, N. Zhang, R. Shao, A. Kaufman, and K. Mueller, "Feature Preserving Distance Fields," *Proc. IEEE Symp. Volume Visualization and Graphics*, pp. 39-46, 2004.
- [27] A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K.D. Harris, "The Asynchronous State in Cortical Circuits," *Science*, vol. 327, no. 5965, pp. 587-590, Jan. 2010.
- [28] L. Sarid, R. Bruno, B. Sakmann, I. Segev, and D. Feldmeyer, "Modeling a Layer 4-to-Layer 2/3 Module of a Single Column in Rat Neocortex: Interweaving in Vitro and in Vivo Experimental Observations," *Proc. Nat'l Academy of Sciences*, vol. 104, no. 41, pp. 16353-16358, Oct. 2007.
- [29] S. Schmitt, J.F. Evers, C. Duch, M. Scholz, and K. Obermayer, "New Methods for the Computer-Assisted 3D Reconstruction of Neurons from Confocal Image Stacks," *NeuroImage*, vol. 23, no. 4, pp. 1283-1298, Dec. 2004.
- [30] J. Schreiner, C. Scheidegger, and C. Silva, "High-Quality Extraction of Isosurfaces from Regular and Irregular Grids," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1205-1212, Sept./Oct. 2006.
- [31] J. Stiles, T. Bartol, and E. Schutter, "Monte Carlo Methods for Simulating Realistic Synaptic Microphysiology Using MCell," *Computational Neuroscience: Realistic Modeling for Experimentalists*, CRC Press, 2001.
- [32] C. Stoll, S. Gumhold, and H.-P. Seidel, "Visualization with Stylized Line Primitives," *Proc. IEEE Visualization (VIS '05)*, pp. 695-702, 2005.
- [33] K. Svoboda and R. Yasuda, "Principles of Two-Photon Excitation Microscopy and Its Applications to Neuroscience," *Neuron*, vol. 50, no. 6, pp. 823-839, June 2006.
- [34] G. Taubin, "A Signal Processing Approach to Fair Surface Design," *Proc. 22nd Ann. Conf. Computer Graphics and Interactive Techniques*, pp. 351-358, 1995.
- [35] J.T. Todd and E. Mingolla, "Perception of Surface Curvature and Direction of Illumination from Patterns of Shading," *J. Experimental Psychology: Human Perception and Performance*, vol. 9, no. 4, pp. 583-595, Aug. 1983.
- [36] B. Ulfhake and S. Cullheim, "A Quantitative Light Microscopic Study of the Dendrites of Cat Spinal Gamma-Motoneurons After Intracellular Staining with Horseradish Peroxidase," *J. Comparative Neurology*, vol. 202, no. 4, pp. 585-596, 1981.
- [37] I. Viola, A. Kanitsar, and M.E. Groller, "Importance-Driven Feature Enhancement in Volume Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 11, no. 4, pp. 408-418, July/Aug. 2005.
- [38] C.C.L. Wang and M.M.F. Yuen, "Freeform Extrusion by Sketched Input," *Computers and Graphics*, vol. 27, no. 2, pp. 255-263, Apr. 2003.
- [39] M.A. Wilson, U.S. Bhalla, J.D. Uhley, and J.M. Bower, "GENESIS: A System for Simulating Neural Networks," *Advances in Neural Information Processing Systems 1*, pp. 485-492, Morgan Kaufmann Publishers, Inc., 1989.
- [40] Z.J. Wood, P. Schroder, D. Breen, and M. Desbrun, "Semi-Regular Mesh Extraction from Volumes," *Proc. Conf. Visualization '00*, pp. 275-282, 2000.
- [41] Y. Zhang, Y. Bazilevs, S. Goswami, C.L. Bajaj, and T.J.R. Hughes, "Patient-Specific Vascular NURBS Modeling for Isogeometric Analysis of Blood Flow," *Computer Methods in Applied Mechanics and Eng.*, vol. 196, nos. 29/30, pp. 2943-2959, May 2007.



Sébastien Lasserre started the PhD degree in 2005, in collaboration with the Blue Brain Project and joined the Laboratoire d'Informatique et de visualisation (LIV) at the Ecole Polytechnique Fédérale de Lausanne (EPFL). His main interests lie in identifying the framework and the languages involved in visualizing large scale scientific models.



Juan Hernando received the bachelor's/master's degree from the Universidad Politécnica de Madrid (UPM) in October 2005. Shortly after, he began the PhD degree at the same university in the field of Visualization. His main interests are in parallel rendering, scientific visualization, and advanced rendering techniques on latest graphics hardware.



Sean Hill received the PhD degree in computational neuroscience from the University of Lausanne in 2000. He is the project manager for Computational Neuroscience of the Blue Brain Project. His research interests include the use of large-scale biologically-realistic computer models to understand information processing, network connectivity and synaptic plasticity in the central nervous system.



Felix Schürmann received the PhD degree in physics from the University of Heidelberg in 2005. He is the general project manager of the Blue Brain Project. The focus of his work was on alternative approaches to computing. Using mixed-signal VLSI he codesigned an efficient implementation of a neural network in hardware and was the first to adopt "Liquid Computing" in hardware.



Pedro de Miguel Anasagasti received the MSc degree in electrical engineering from Carnegie-Mellon University and the PhD degree in electrical engineering from the Universidad Politécnica de Madrid. He is a full professor in Computer Science at the UPM, Spain and former director of the Centro de Supercomputación y Visualización de Madrid. His research interests include high performance computing and visualization.



Henry Markram is the director of the Blue Brain Project. He moved to the EPFL in 2002 as a full professor and founder/director of the Brain Mind Institute (BMI). At the BMI, he has continued to unravel the blueprint of the neocortical column, building state-of-the-art tools to carry out multi-neuron patch clamp recordings combined with laser and electrical stimulations.



Georges Abou Jaoudé drives the LIV at the School of Architecture at the EPFL. Assistant to Professor Charles Rapin from 1984 to 1988, he became full professor of Architecture and Computer Science at the EPFL in 1992. His teaching covers the modeling, representation, and simulation domains.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**