Medical Image Analysis 13 (2009) 167–179

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Robust 3D reconstruction and identification of dendritic spines from optical microscopy imaging

Firdaus Janoos [a,b], Kishore Mosaliganti [a,c], Xiaoyin Xu [d,e], Raghu Machiraju [a,c], Kun Huang [c], Stephen T.C. Wong [f,g,*]

[a] Department of Computer Science and Engineering, The Ohio State University, Columbus, OH, USA
[b] Davis Heart and Lung Research Institute, The Ohio State University, Columbus, OH, USA
[c] Department of Biomedical Informatics, The Ohio State University, Columbus, OH, USA
[d] Harvard NeuroDiscovery Center, Harvard Medical School, Boston, MA, USA
[e] Functional and Molecular Imaging Center, Department of Radiology, Brigham and Women's Hospital, Boston, MA, USA
[f] Research Division and Medical Physics Section, Department of Radiology, The Methodist Hospital, Houston, TX, USA
[g] Center for Biotechnology and Informatics (CBI), The Methodist Hospital Research Institute, Houston, TX, USA

ARTICLE INFO

Article history:
Received 16 April 2007
Received in revised form 21 April 2008
Accepted 23 June 2008
Available online 24 July 2008

Keywords:
Neuron
Dendrite
Spine
3D reconstruction
Curve-skeleton
Medial geodesic function

ABSTRACT

In neurobiology, the 3D reconstruction of neurons followed by the identification of dendritic spines is essential for studying neuronal morphology, function and biophysical properties. Most existing methods suffer from problems of low reliability, poor accuracy and require much user interaction. In this paper, we present a method to reconstruct dendrites using a surface representation of the neuron. The skeleton of the dendrite is extracted by a procedure based on the medial geodesic function that is robust and topology preserving, and it is used to accurately identify spines. The sensitivity of the algorithm on the various parameters is explored in detail and the method is shown to be robust.

1. Introduction

In neurobiology, 3D reconstruction of neuronal structures such as dendrites and spines is essential for understanding the relationship between their morphology and functions (Dima et al., 2001) and helps us understand neuronal circuitry and behaviour in neurodegenerative diseases. Dendritic spines, or spines in short, play a significant role in many neurological conditions (Kasai et al., 2003; Nimchinsky et al., 2002). These conditions include Alzheimer's disease (Spires et al., 2005), Parkinson's disease (Comery et al., 1997), tuberous sclerosis complex syndrome (Tavazoie et al., 2005), and other neurological diseases. Dendrites are tree-like structure of a neuronal cell and spines are small protrusions formed on the surface of a dendrite. Spines can assume different shapes and appear and disappear totally (Hering and Sheng, 2001) and their morphological changes have been proven to be associated with synaptic plasticity (Yuste and Bonhoeffer, 2001). Important aspects of cognitive function, such as experience-based learning (Engert and Bonhoeffer, 1999), attention and memory (Moser et al., 1994) are correlated with variations in dendritic arborescence and with spine density and distribution (Glantz and Lewis, 2000; Zito et al., 2004). It is hypothesized that the dendritic spine structure affects the physiological properties of synapses located on them (Harris et al., 1992; Benshalom and White, 1988; Harris and Stevens, 1988). The topology (branching structure) of neuronal dendrites and their spines underlie the connectivity of neural networks and may therefore be important predictors of their function (Koh et al., 2002) (see Fig. 1).

3D light microscopy images of neurons have a significant amount of information that is lost when projecting to 2D. For example, structures that are orthogonal to the imaging plane, or structures that overlap each other along the imaging axis cannot be identified. See Fig. 2 for an illustration of this. The preferred method for detailed study of cell morphology and topology is from 3D reconstructions (Stevens and Trogadis, 1984; Wilson et al., 1987). Automatic reconstruction aids in the analysis of a large number of neurons and the mapping of the spatial relationships between different tracts and neuropiles, such that detailed

* Corresponding author. Address: Research Division and Medical Physics Section, Department of Radiology, The Methodist Hospital, Houston, TX, USA
E-mail addresses: janoos@cse.ohio-state.edu (F. Janoos), mosaligk@cse.ohio-state. edu (K. Mosaliganti), xxu@bwh.harvard.edu (X. Xu), raghu@cse.ohio-state.edu (R. Machiraju), khuang@bmi.osu.edu (K. Huang), stwong@tmhs.org (S.T.C. Wong).

1361-8415/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.media.2008.06.019

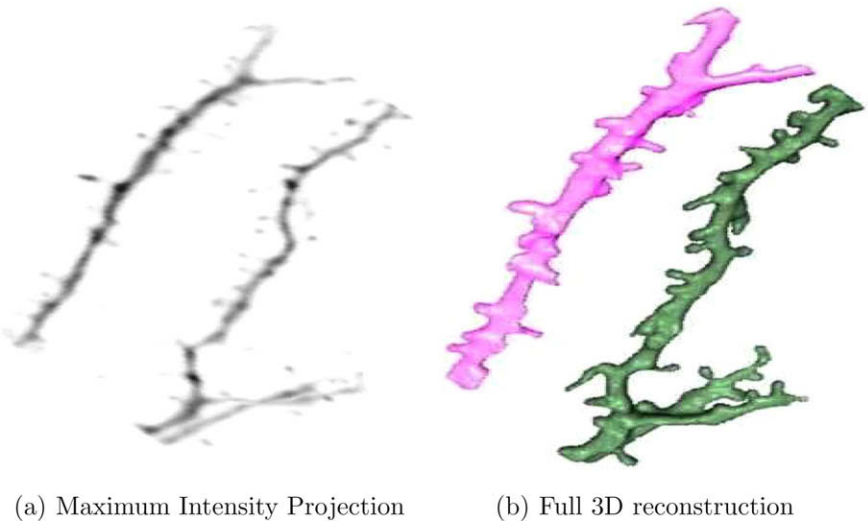

Fig. 1. (a) 2D maximum intensity projection of a 3D neuron image, in which information along the imaging axis is lost. (b) Full 3D reconstruction. The 3D reconstruction captures the branching structure and the morphology of the dendrites with much greater accuracy, and is used as the basis for spine identification.

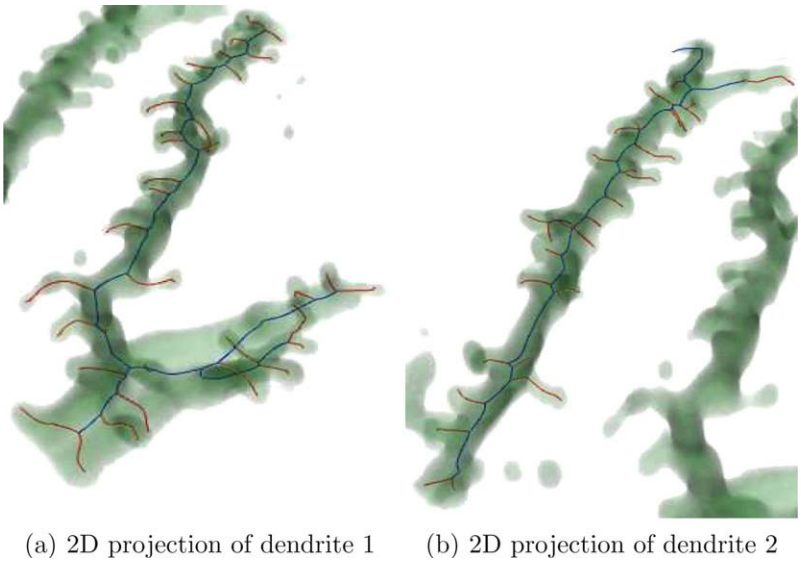

Fig. 2. (a) and (b) 2D projections of the neuron images, with the backbone (blue) and spines (red) identified by the algorithm overlaid on them. From these images, it can be seen that in 2D a large amount of the neuron structure is lost.

mathematical and physical models can be constructed to estimate those physiological parameters which can not be otherwise measured easily (Dima, 2002). Graph models (*dendrograms*) of the neuron, which concisely capture the neuron geometry and topology, are extremely valuable for analyzing the structure of the neuron backbone and the dendritic spines (Weaver et al., 2004) and elucidate their synapses and neurological functions (Zito et al., 2004). As discussed above, many neuronal functions are observed to be correlated with the appearance or disappearance of neuronal structures and the morphology of the spines. It is therefore important to develop robust 3D reconstruction method to trace the dendrites and detect the dendritic spines on them. There has been extensive research in the field of neuron reconstruction, dendrite tracing and spine identification. However, most of these methods suffer from low reliability and poor accuracy and have heavy requirements of manual supervision.

In a related work (Mosaliganti et al., 2006), we developed a method to perform automatic temporal tracking and matching of spine evolution. Our method, however, suffered from poor spine detection, and the inadequacies of existing methods to meet these needs has motivated our current work on developing robust algorithms for spine identification and morphometry, whereas existing methods for neuron reconstruction treat the neuron as a volume, we choose to represent the dendrite and the spines as a surface (2-manifold). The advantages of a surface representation are that it lets us enforce physically plausible smoothness constraints on the shape of the neuron, and it facilitates the extraction of the neuronal skeleton. There are a large number of algorithms and heuristics in literature for skeletonisation, each of which produce different results and do not provide guarantees of correctness (Cornea et al., 2005; Ma et al., 2003). We use a geometric skeletonisation method by Dey and Sun (2006b) based on the *medial geodesic function*. This algorithm has been proved to be well-posed and robust against noise, and produces curve-skeletons that preserve the topology of the original object (homotopic equivalence). The surface model of the dendrite along with the information computed by the skeletonising procedure allows for the accurate identification of spines, and for morphological measurements like

the spine diameter and eccentricity at the base and tip, spine length, volume, etc. The parameters in the spine identification procedure are intuitive, easy to understand, robust, and are motivated by the biology of the neuron being studied.

This paper is organized as follows: in Section 2, we review some important contributions in neuron reconstruction, tracing, dendrogram construction and spine identification. Next, we describe the data-set, image acquisition and processing methods in Section 3 (see Fig. 3a). The surface extraction, neuron reconstruction, skeletonisation and spine identification pipeline is explained in Section 4 (see Fig. 3b). In Section 5 we demonstrate the results of our method applied on the data-set, and quantitatively evaluate the accuracy of spine detection and morphometry. We also explore the implications of the various parameters in detail and make recommendations for selecting their values. Finally, in Section 6 we conclude by discussing some thoughts on our current method and directions of future work.

2. Related work

Most of the current methods of reconstruction are semi-automatic and require user guidance to identify salient structures and deal with ambiguities in each confocal microscopy data-set independently (Glaser and Glaser, 1990; Garvey et al., 1973; Carlbom et al., 1994). These techniques demand several weeks of a specialist's time for one neuronal reconstruction, and do not have the objectivity of automated methods. The goal of decreasing the expense of user interaction often acts contrary to that of ensuring the accuracy and the topological correctness of the result. Consequently, a few recent methods (e.g. Schmitt et al., 2004) seek for an optimal compromise been automatic segmentation and manual reconstruction.

Some reconstruction methods represent neuronal (and other tubular branching) structures as a tree of connected cylinders or similar mathematical objects (Urban et al., 2006; Al-Kofahi et al., 2002; Uehara et al., 2004; Tyrrell et al., 2006; Schmitt et al., 2001; Herzog et al., 1997; Herzog et al., 1998), which detracts from the capability of the model to provide accurate quantitative measurement of the underlying cellular morphology, and require strong assumptions about the objects of interest. Due to the morphologic complexity and variety of neuronal cell types, no general models are suggested in literature.

Other methods (Cohen et al., 1994; Goldbaum et al., 1990a; Goldbaum et al., 1990b; Clark et al., 1992; Chaudhuri et al., 1989; Dima et al., 2001; He et al., 2003; Can et al., 1999; Gerig et al., 1993) employ curve-skeleton methods to build a linear graph representation of dendritic structures. These methods conform to a weak model, which implies that only very weak assumptions are made about the shape of the objects of interest. However, the curve-skeletonisation algorithms employed tend to be very sensitive to small changes in the object shape, and moreover do not provide topological guarantees about the resulting 1D skeleton. We discuss this some more in Section 4.2.

The analysis of dendritic structure and morphology is largely accomplished manually and is extremely time consuming, not reproducible, and its accuracy is dependent on the skill of the user. A few spine identification and quantification techniques of varying degrees of automation have been suggested to reduce manual labour and improve the accuracy and reproducibility of the result, none of which has apparently been used and verified on large data-sets. Some authors (Rusakov and Stewart, 1995; Weaver et al., 2004; Mosaliganti et al., 2006) use the medial axis to identify spines in 2D as protrusions relative to dendritic skeleton. We have noticed from our earlier work that there is significant information in the 3D image which is lost when projecting to 2D, and the accuracy of such methods is limited. Others (Watzel et al., 1995; Koh, 2001) use 3D medial axis-based strategies to extract a skeleton, and identify spurs as potential spines. Medial axis-based methods suffer from the general problems of sensitivity to noise and spurious spine detections, and have to use heuristics to eliminate false positives. Model based spine identification techniques (Herzog et al., 1998; Al-Kofahi et al., 2002) have trouble detecting thin-necked and short spines and have to be manually supervised. In Koh et al. (2002), the authors devise a 3D technique in which spines are not detected using the medial axis branches emerging from the backbone, but instead, as geometric protrusions relative to the backbone. The method contains several parameter settings that require extensive fine tuning.

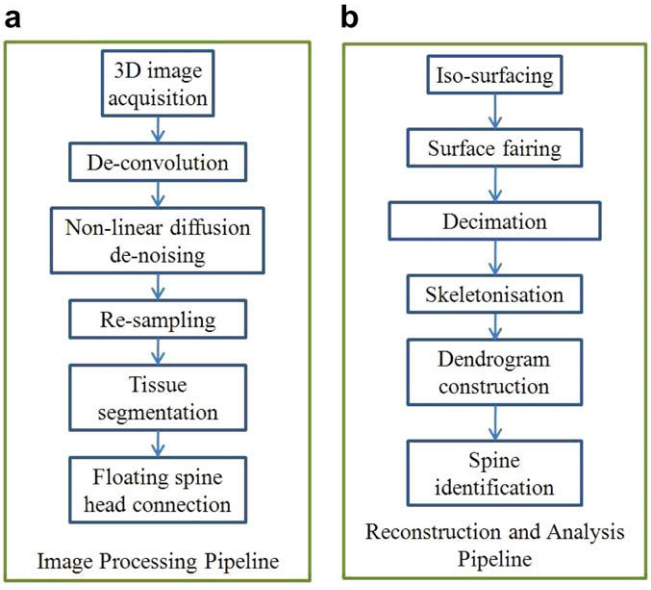

Fig. 3. Algorithm overview: (a) image processing pipeline and (b) reconstruction and analysis pipeline.

3. Image acquisition and processing

In this section we describe the data-set (Section 3.1), followed by the image processing pipeline that corrects for the anisotropy in the sampling resolution (Section 3.2) and segments out the neuronal cytoplasm from the background phase (Section 3.3). This is followed by a step to join the floating spine heads and to remove spurious tissue fragments in the sample (Section 3.4).

3.1. Data-set

3D images of pyramidal neurons in rat hippocampi expressing GFP were acquired by the digitization of neuronal cultures using a two-photon laser scanning microscopy with a 40× objective and 0.8 NA (Zito et al., 2004). The image stacks have dimension $512 \times 512 \times 12$ voxels at $0.07\ \mu m \times 0.07\ \mu m \times 1\ \mu m$ resolution. For high content neuron screening green fluorescent protein (GFP) is used to mark neurons *in vitro*. GFP absorbs blue light and converts it to green light which is of lower energy. The emitted green light is then captured by an optical microscope such as a two-photon laser scanning microscope (2PLSM).

To correct the images for the microscope's point spread function (PSF), which causes out-of-focus objects to appear in the optical slices, we use the de-convolution package AutoDeblur[1] to restore

[1] AutoDeblur is a product of AutoQuant Image Inc.

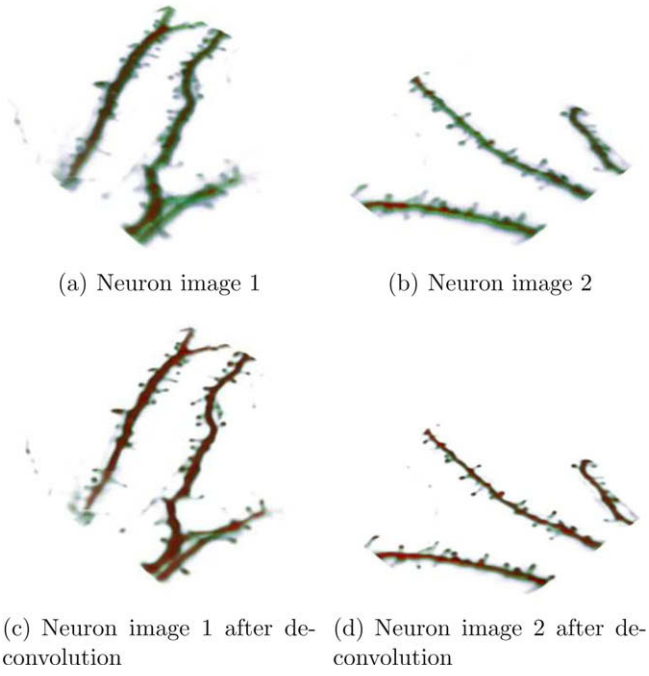

Fig. 4. Volumetric renderings of GFP stained pyramidal neurons from rat hippocampi before (a) and (b) and after (c) and (d) de-convolution. The data-set resolution is $512 \times 512 \times 12$ at 1:1:10 voxel aspect ratio ($0.07\ \mu m \times 0.07\ \mu m \times 1\ \mu m$).

the image. Fig. 4 shows raw and deblurred images. The deblurred image was obtained after 10 iterations. The parameters of the de-convolution algorithm depend on the setup of the microscopy.

The nature of the image acquisition process and photo-bleaching effects introduces the following types of artifacts: (i) photon shot noise; (ii) presence of unrelated structures; and (iii) floating spine heads. These artifacts require the additional processing steps described next.

3.2. De-noising and re-sampling

The intensity profile in neuronal regions exhibits noise and sharp variations in contrast, especially along the cell boundaries. Fig. 5a and b shows a small portion of an $X - Y$ slice from a 3D neuron image, and the corresponding *Monge* map (intensity map). To de-noise the image, non-linear diffusion filtering (Perona and Malik, 1990) is used, which removes high-frequency noise while avoiding the blurring and localization problems of linear Gaussian smoothing (Witkin, 1983). If Ω denotes the domain of the image $g(\mathbf{x}) : \Omega \to \mathbb{R}$ then the filtered image $u(\mathbf{x}, t)$ is the solution to the non-linear diffusion equation

$$\frac{\partial u}{\partial t} = \operatorname{div}(\mathbf{D}(|\nabla u|^2)\nabla u) \quad \text{on } \Omega \times (0, \infty), \quad \text{where} \tag{1}$$

$$\mathbf{D}(|\nabla u|^2) = \frac{1}{1 + |\nabla u|^2/\lambda_{nl}^2} \quad (\lambda_{nl} > 0) \tag{2}$$

with the original image as the initial condition $u(\mathbf{x}, 0) = g(\mathbf{x})$ on Ω, and Neumann boundary conditions $\partial_n u = 0$ on $\partial\Omega \times (0, \infty)$.[2] Here λ_{nl} plays the role of a contrast parameter (Weickert, 1996) smoothing low contrast areas ($|\nabla u| \leqslant \lambda_{nl}$) and enhancing high contrast areas ($|\nabla u| > \lambda_{nl}$).

Unfortunately, this diffusion equation is numerically unstable and Catté et al. (1992) propose a regularization by convolving u with a Gaussian K_σ to compute the gradient ∇u_σ. To account for the $1:1:10$ anisotropy in the imaging resolution, we modify the Catté filter to use an anisotropic Gaussian kernel K_Σ with a $10:10:1$ ratio of scales $\sigma_x, \sigma_y, \sigma_z$. The non-linear diffusion equation with anisotropic regularization has the form

$$\frac{\partial u}{\partial t} = \operatorname{div}(\mathbf{D}(|\nabla u_\Sigma|^2)\nabla u_\Sigma) \quad \text{on } \Omega \times (0, \infty), \tag{3}$$

$$\text{where} \quad \Sigma = \sigma\ \operatorname{diag}\left(1, 1, \frac{1}{10}\right),$$

$$u_\Sigma = K_\Sigma \bigstar u. \tag{4}$$

Here, $\bigstar$ represents convolution. Fig. 5c and d shows a 2D section of the results of (3D) non-linear filtering. For comparison, the results of standard Gaussian filtering (with anisotropic scales) are shown in Fig. 5e and f. We see that the non-linear diffusion method has effectively removed noise, while preserving tissue boundaries.

After smoothing, the image is re-sampled to 1:1:1 resolution by down-sampling by three in the X and Y, while up-sampling by a factor of 3.5 in the Z directions. We found that quartic (fourth order) B-spline interpolation provides the best results, by visually examining the Monge maps before and after interpolation. Fig. 6a and b shows the results at the end of the de-noising and re-sampling stage.

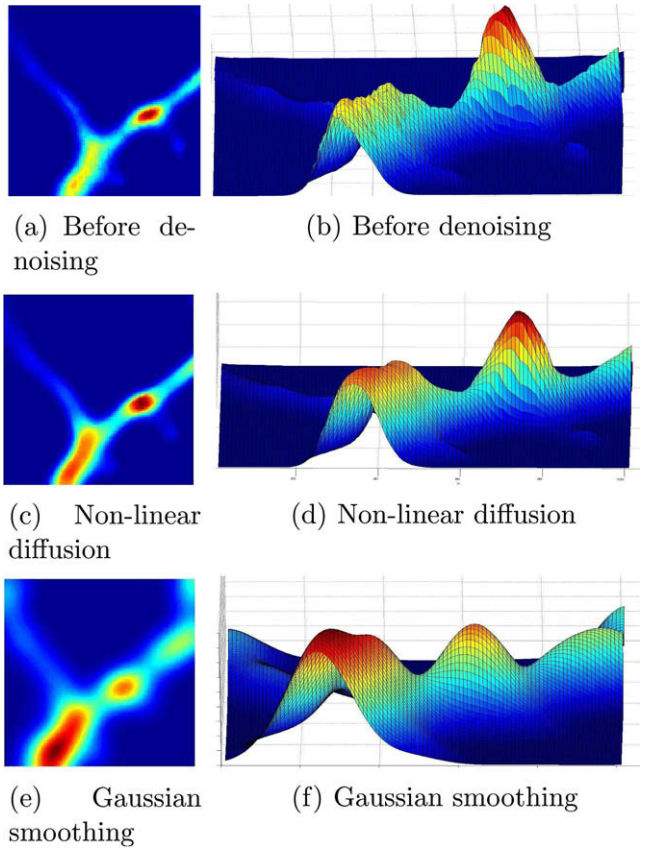

Fig. 5. Results of non-linear diffusion de-noising with anisotropic regularization. The left column shows a section of 2D slice of the 3D neuron image, and the right shows the corresponding Monge map. (a) and (b) Image before smoothing, (c) and (d) results for non-linear diffusion, while (e) and (f) result of Gaussian smoothing. We can see that the non-linear diffusion method de-noises the image while preserving edges. Both Gaussian and non-linear filtering were carried out in 3D.

3.3. Segmenting the neuron

Because of the edge enhancing property of non-linear diffusion, global thresholding of the intensity field gives a reliable

[2] ∂_n denotes the derivative normal to the image boundary.

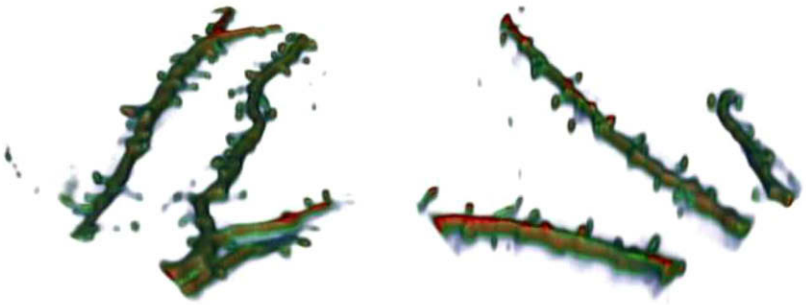

(a) Neuron image 1 after de-noising and re-sampling (b) Neuron image 2 after de-noising and re-sampling

Fig. 6. Neuron volumes after de-noising and re-sampling (Section 3.2) to 1:1:1 voxel aspect ratio.

and accurate segmentation of the neuron. This is because the variations in the intensity field outside the neuron become almost zero and hence any threshold Ω above zero is very effective suppressing the artifacts from random intensity fluctuations in the background image. At the same time, the filter sharpens edges at the boundaries of the neuronal objects thereby making the segmentation more robust to variations in the threshold parameter Ω. The conditioning of the thresholding operation is quantified by the relative change in the size of the segmented object with respect to a relative change in Ω. Assuming that the neurons are approximately cylindrical objects, the condition number $\mathrm{cond}(\Omega)$ is given by

$$\mathrm{cond}(\Omega) = \frac{\delta r/r}{\delta\Omega/\Omega}, \tag{5}$$

where r is the radius of the cylinder. The value of $\delta r/r$ is approximated by

$$\frac{\delta r}{r} \approx \sqrt{\frac{\delta V}{V}}, \tag{6}$$

where V is the volume of the segmented dendritic tissue. We present our investigation into the sensitivity of Ω in Section 5.2. Fig. 7 shows volumetric renderings of segmented neuronal objects from two data-sets.

3.4. Floating spine heads

The neuron images contain floating spine heads, separated from the main dendritic backbone, due to photo-bleaching effects and limited microscope resolution. The neuron sample also contains unrelated tissue fragments that show up as disconnected blobs during imaging (Fig. 7). While the spurious fragments should be removed, the floating spine heads must be identified and connected back to the main dendrite. An important consideration when making this join is that the true topology of the dendrite be recreated as best as possible. At the same time, approximating the morphology of such spines is acceptable, since this information is fundamentally not present in the acquired images.

It has been observed that the floating spine heads tend to point towards the dendrite, while the dendrite too has a protrusion pointing towards the spine, and both are in close proximity of each other. Fig. 8a shows a 2D projection of a segmented neuron image with three detached spine heads (circled in red) and a few spurious fragments. For spine heads 1 and 3 locating the closest point on the main backbone is sufficient to determine the connection between the two. However, for case 2, a join to the closest point (yellow arrow) is incorrect. Here, along with the orientation of the spine head itself, the protuberance in the dendrite (green arrow) must used to guide the connection. This suggests a method of connecting floating spine heads by growing the spine and the dendrite along the directions in which they protrude. If a join occurs within a certain number of time-steps it is a valid spine, else it is spurious fragment.

Active contour shape models (Malladi et al., 1995) are a level-set approach that enables us to achieve these goals, by using shape information to control the evolution of iso-contours. The boundary $\zeta(t) = \{\mathbf{x}|\psi(\mathbf{x},t) = 0\}$ of the neuron at time t is encoded as the zero level-set of the field $\psi(\mathbf{x},t)$. The update equation of the level-set is given by

$$\frac{\partial\psi}{\partial t} = g(\psi)(c+\kappa)|\nabla\psi|, \quad \text{where } g(\psi) = \frac{1}{1+|\psi|^2}. \tag{7}$$

The parameter κ is the curvature along the normal to the level-set contour and c is a balloon force that evolves the level-set front outwards. The term $\kappa|\nabla\psi|$ guides the evolution of the front using the local curvature.

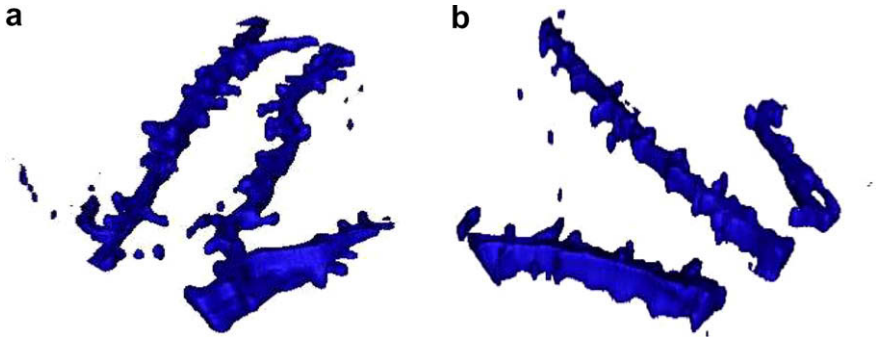

Fig. 7. Volumetric rendering of segmented neuronal objects: (a) neuron image 1 after segmenting and (b) neuron image 2 after segmenting.

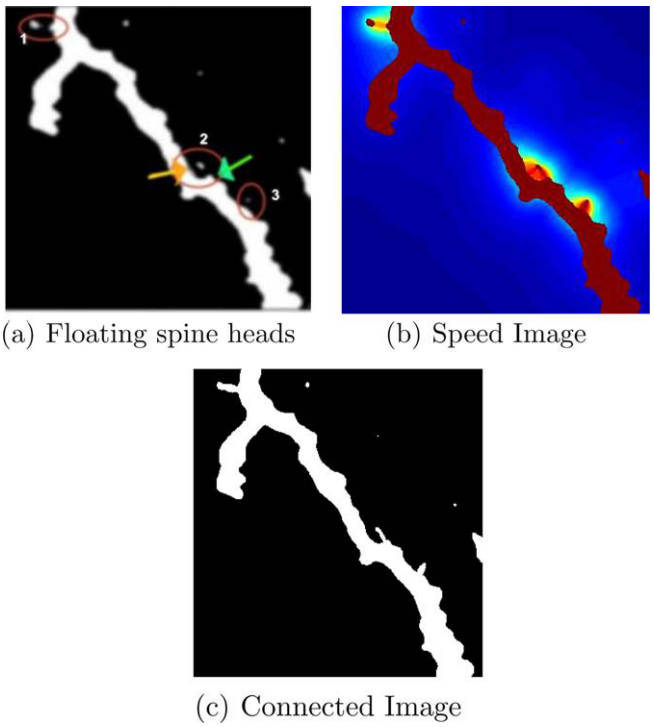

Fig. 8. Illustrative example (in 2D) of the floating spine head problem and its solution: (a) neuron image after segmentation with floating spine heads and spurious tissue fragments in 3D; (b) speed function that controls the level-set evolution; and (c) floating spine head connected to the backbone.

Therefore, the complete algorithm to connect floating spine heads and discard fragments is as follows:

(i) From the segmented image I, identify dendrites and potential floating spines by thresholding the volume of the connected components with respect to a reference volume ϑ (see Section 5.2 for more details).
(ii) Set the signed distance map $D(I)$ of the neuron image as the level-set function $\psi(\mathbf{x}, 0)$.
(iii) Evolve the level-set (Eq. (7)) to move in the directions of high curvature for a specified number of time-steps τ_{gac} (Fig. 8b).
(iv) If the level-sets originating from spine heads and dendritic backbones meet within τ_{gac}, then determine the location of their join point and connect them up at that point (Fig. 8c).
(v) Tag all remaining fragments as spurious, and discard them.

Because of the proximity constraint, the level-set needs to be evolved only over a very few number of time-steps ($\tau_{\text{gac}} \approx 5$) for a valid join to occur. The distance transform, too, needs to be computed only in the close vicinity of the edges and can be done very efficiently in linear time (Breu et al., 1997). Therefore, the overall computation cost is of this procedure is relatively low.

4. Neuron reconstruction and analysis

After the image processing pipeline, the connected components are identified as separate dendrites and are analyzed independently. Existing neuron reconstruction methods use signal processing and image processing techniques, followed either by model fitting or skeletonisation for neuron reconstruction (Section 2). We represent the dendrite by a surface model because (a) it allows us to control the smoothness of the neuron surface thereby imposing a physically plausible constraint on the reconstruction and (b) it lets us perform detailed morphological measurements.

4.1. Surfacing and surface fairing

The surface of the segmented dendrite is obtained by iso-surfacing at any value between (0,1), using extended marching cubes (Nielson and Hamann, 1991). This surface is over tesselated and suffers from artifacts of staircase noise (Fig. 9a). In signal processing terms, we need to first low-pass filter the surface to remove the high-frequency noise, and then down-sample it to a sufficient resolution.

Low-pass filtering is effected using the two-step surface fairing method described by Taubin (1995). Let $\mathbf{x} = (x_1, x_2, x_3)^{\mathrm{T}}$ be the 3D coordinates defined at the vertices of a polyhedral surface. The Laplacian of the a point $\mathbf{x}_i$ on the surface is defined by the weighted average over its neighbourhood $\mathbb{N}_i$ as

$$\Delta \mathbf{x}_i = \frac{1}{|\mathbb{N}_i|} \sum_{j \in \mathbb{N}_i} w_{ij} (\mathbf{x}_i - \mathbf{x}_j). \tag{8}$$

The fairing process is a smoothing step that also causes surface shrinkage (Eq. (9)) followed by an expansion step (Eq. (10)), applied iteratively N times:

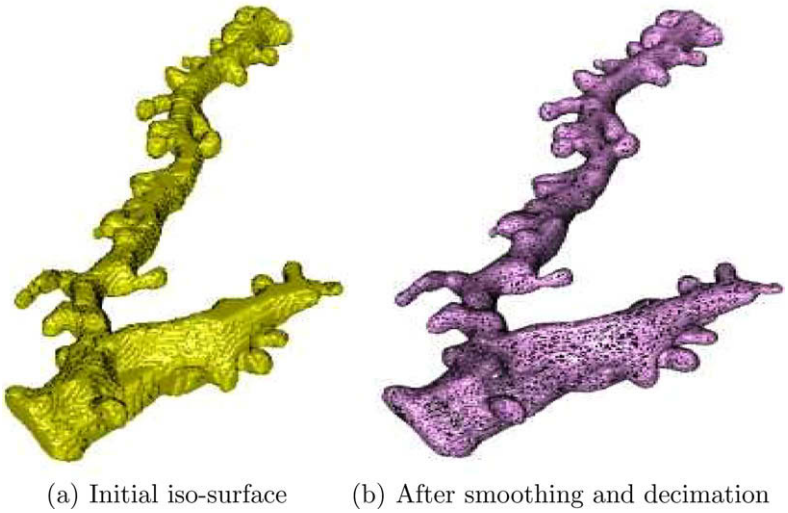

Fig. 9. Surface model of dendrite: (a) original iso-surface of the segmented neuron object and (b) surface after low-pass filtering and quadric error decimation.

$$\mathbf{x}'_i = \mathbf{x}_i + \alpha\Delta\mathbf{x}_i \quad \text{for } 0 < \alpha < 1, \tag{9}$$

$$\mathbf{x}''_i = \mathbf{x}'_i + \mu\Delta\mathbf{x}_i, \quad \text{where } \mu < -\alpha. \tag{10}$$

The transfer function $f(k)$ of the filter, with respect to surface frequency k, has the following property:

$$\lim_{N\to\infty} f^{(N)}(k) = \begin{cases} 1 & 0 \leqslant k \leqslant \frac{1}{\alpha} + \frac{1}{\mu}, \\ 0 & \frac{1}{\alpha} + \frac{1}{\mu} < k \leqslant 2. \end{cases} \tag{11}$$

Here $(1/\alpha + 1/\mu) > 0$ is the cut-off frequency and N the number of iterations, controls the rate of decrease in the stop-band. This algorithm is fast (linear time), produces smoothing without shrinkage, and quickly achieves a stable solution with respect to N.

Next, the tessellation density is reduced by decimating the mesh using the quadric error metric (Garland and Heckbert, 1998). Here, every edge is assigned a cost function, namely the error resulting from its contraction, and the lowest cost edges are iteratively selected and collapsed. Each vertex is associated with a set of planes, and the error at the vertex is defined to be the sum of squared distances from it to all the planes in its set. Each set is initialized with the faces incident to the vertex in the original surface. When an edge is contracted into a single vertex, the resulting set is the union of the two sets associated with the endpoints. The cost of contracting an edge (v_1, v_2) to a single vertex $\bar{v}$ is now the error at $\bar{v}$. This decimation technique does not prevent changes of topology in the mesh, and it suffers from small inaccuracies. However, given the simple topology of a dendrite, and the speed and simplicity of the algorithm, this method is very appropriate for our application. Fig. 9b shows the result of this decimation step.

Increasing the surface decimation factor ρ has two benefits: (i) it simplifies the model and makes further computations more efficient and (ii) it allows us to impose smoothness constraints on the model. The smoothness of the surface is measured by the average dihedral angle ($\overline{\phi}_{\mathbb{M}}$) of the edges in the mesh. This is further elaborated upon in Section 5.2.

4.2. Curve-skeletonisation

A curve-skeleton is a 1D curve, possibly with branches, in the 'center' of the shape. A related and much more well-defined concept is the medial axis which is also referred to as the skeleton. For a 3D shape, however, the medial axis has two-dimensional components (*medial surface*). Therefore, the medial axis cannot be a substitute for a 1D skeleton. Another disadvantage of the medial surface (axis) is its intrinsic sensitivity to small changes in the object's surface (Choi and Seidel, 2002). Essentially, any curve-skeleton should satisfy the following basic properties (Cornea et al., 2005):

(1) Homotopic to the original object (topology preservation).
(2) Invariant under isometric transformations.
(3) Allow recovery of the original object (reconstruction).
(4) 1D (thin).
(5) Centered within the object.
(6) Visibility of every boundary point on the object from at least one curve-skeleton location (reliable).
(7) Ability to distinguish different components of the original object, reflecting its part/component structure (junction detection).
(8) Preservation of the connectedness of the original object.
(9) Small changes in the skeleton for small changes in object surface (robust).

We use the definition of curve-skeletons based on the *medial geodesic function* by Dey and Sun (2006b), which combines the intrinsic property of the surface (the geodesic distances) along with its embedding in $\mathbb{R}^3$ (the medial axis) thereby capturing the shape information comprehensively. The medial geodesic function gives the shortest geodesic distances between the points, where the maximal balls centered at the medial axis touch the surface. Formally, if $O \subset \mathbb{R}^3$ is a space called *shape* bounded by a connected *2-manifold* surface S, then the medial axis $M \subset O$ is the set of centers of the maximal balls inscribed in O. Let $M_2 \subset M$ be the set of points on the medial axis whose maximal balls touch the surface S at two distinct points. It can be shown that M_2 is also a 2-manifold and covers most of M (i.e. $M \setminus M_2$ has measure 0). For a point $x \in M_2$, let B_x be the maximal inscribed ball centered at x and a_x and b_x be the two touching points where B_x meets S. Then $f(x)$, the length of the geodesic path on S between a_x and b_x, is the medial geodesic function (MGF). The curve-skeleton is defined as the singular set (maxima or saddle points) of $f(x)$ for $x \in M_2$.

It is mathematically shown that this definition of a curve-skeleton has properties of homotopic equivalence, isometric invariance, thinness (1D), centeredness, junction detection, stability (robustness) and connectedness. The MGF values at each point on the skeleton give the size information of the shape. Also the ratio ϵ between the geodesic and Euclidean circles passing through the touching points quantify how different the shape is from a tubular one (*eccentricity*).

The algorithm has one parameter $-1.1 \leqslant \theta \leqslant 0.0$ that controls the strictness for selecting points from the medial axis M_2 as being skeleton points. As θ decreases the curve skeleton becomes less detailed. Formally, if SK_S^θ is the curve-skeleton for surface S extracted with parameter θ, then $SK_S^{\theta_1} \subseteq SK_S^{\theta_2}$ if $\theta_1 < \theta_2$. The selection of this parameter is explained in Section 5.2. Fig. 10a–c shows the results of curve-skeleton extraction for a few values of θ.

4.3. Dendrite tree model

The curve-skeleton SK_S^θ of the neuron is represented by an attributed tree structure $\mathscr{D} \equiv \{\mathscr{V}, \mathscr{E}\}$ which compactly encodes the geometry and topology of the dendrite. The set of vertices is

$$\mathscr{V} = \{v_i | v \equiv (\mathbf{x}, d)\}, \tag{12}$$

where $\mathbf{x} \in \mathbb{R}^3$ are the spatial coordinates and $d \in \mathbb{N}$ is the degree of the vertex. The set of edges is

$$\mathscr{E} = \{e | e \equiv ([v_i, v_j], \gamma[\mathbf{a}, \mathbf{b}], r, f, \epsilon)\}, \tag{13}$$

where $v_i, v_j \in \mathscr{V}$ are its vertices, γ is the length of the edge, $[\mathbf{a}, \mathbf{b}] \in \mathbb{R}^3$ are the touching points of medial ball with the surface, r is the radius of the medial ball, $f \in \mathbb{R}$ is the medial geodesic length for the edge, and ϵ is the eccentricity.

The dendrite tree has two types of chains of edges $\{e_1(v_1, v_2), e_2(v_2, v_3), \ldots, e_n(v_n, v_{n+1})\}$:

(1) Branch chains that start at leaf node $(v_1 | d(v_1) = 1)$ and end at a branch node $(v_{n+1} | d(v_{n+1}) > 2)$.
(2) Backbone chains that run between two branch nodes.

Fig. 10d shows the graph of the dendrite in Fig. 10b. The dendrite tree is pruned of those branch chains whose cumulative length $\sum_{i=1}^{n} \gamma_i$ is less than a threshold length $\Gamma_{\min}$. This step eliminates spurious branches in the curve-skeleton by imposing restrictions on the minimum length of potential spines in the dendrite.

4.4. Morphometry and spine identification

For each branch β in the dendrogram model we compute the following morphological features:

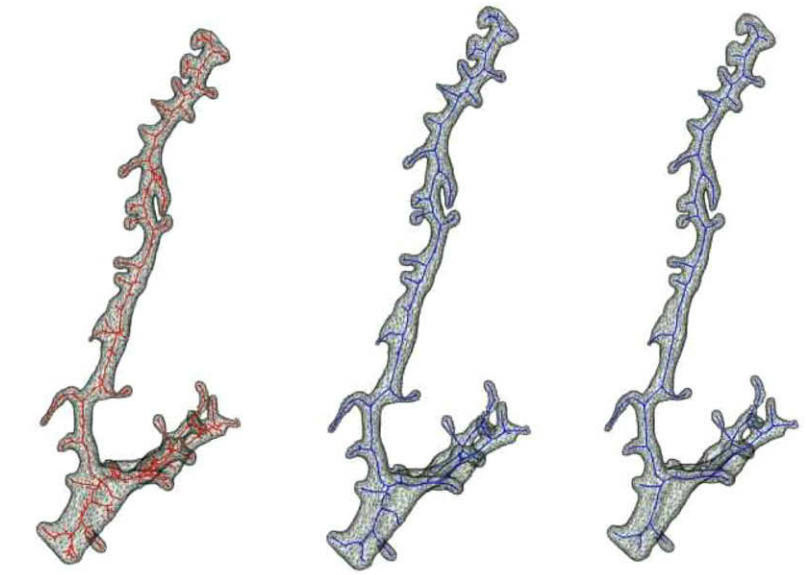

(a) SK_S^{θ} for $\theta = 0.0$ (b) SK_S^{θ} for $\theta = -0.2$ (c) SK_S^{θ} for $\theta = -0.5$

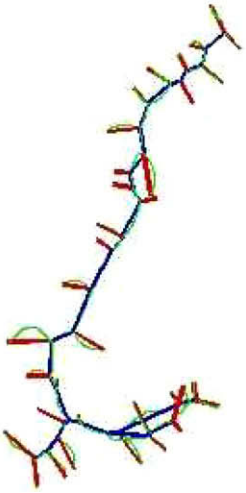

(d) Dendrite graph for $SK_S^{(-0.2)}$

Fig. 10. (a)–(c) Curve-skeletons for $\theta = 0.0, -0.2, -0.5$, respectively; (d) dendrite graph (before pruning) for the skeleton in (b). The green curve is the skeleton, the blue lines indicate backbone edge chains, and red lines are branch chains.

- length γ^{β} as the cumulative lengths $\sum_{i=1}^{n}\gamma_i$ of the edges $e_1, \ldots, e_n$ from the base (branch nodes) to the apex (leaf nodes),
- radius r_0^{β} of the medial ball and eccentricity ϵ_1^{β} at the base of the branch,
- maximum radius $r_{\max}^{\beta} = \max_{i=1}^{n} r_i$,
- minimum radius $r_{\min}^{\beta} = \min_{i=1}^{n} r_i$,
- average weighted radius $r^{\beta} = (1/\gamma_i^{\beta})\sum_{i=1}^{n}\gamma_i r_i$,
- volume V^{β} as the cumulative volume of the edges $\sum_{i=1}^{n}\pi r_i^2\gamma_i$ (approximating each edge by a cylinder),
- the angle ϕ^{β} the branch chain makes with the backbone edge at the base,
- the curvature κ^{β} of the backbone at the base of the branch chain.

To distinguishing between branch chains that belong to end segments of the dendritic backbone and those that belong to the spines (see Fig. 10d) in the pruned dendrite tree, we use the following decision sequence:

(1) If β is a branch chain of cumulative length γ^{β} greater than a threshold length $\Gamma_{\max}$, then mark it as dendritic backbone.
(2) Else:
  (2i) If the average weighted radius r^{β} is greater than a threshold $R_{\max}$, mark it as a backbone chain β_{bb}.
  (2ii) Else mark it as spine chain β_{s}.

The threshold $\Gamma_{\max}$ enforces a maximum length constraint on valid spines, while $R_{\max}$ enforces a maximum radius constraint. While this framework allows for more sophisticated morphological

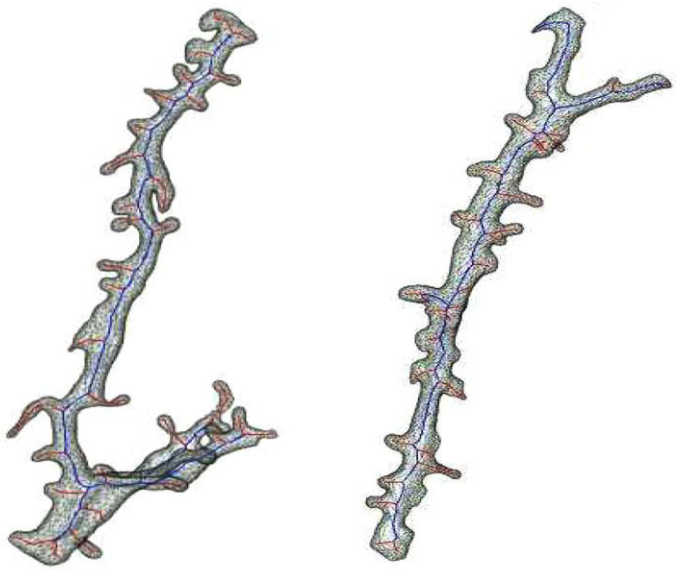

(a) 3D visualization of dendrite 1

(b) 3D visualization of dendrite 2

Fig. 11. (a) and (b) 3D visualizations of two dendrites, their backbones (blue) and spines (red).

models when checking for spines, we find that the above model performs well for our data-sets.

5. Results

Our method was implemented in a combination of MATLAB®,[3] C++, ITK[4] and VTK.[5] We used the CurveSkel software package (Dey and Sun, 2006a) to calculate the curve-skeleton of the dendrites. Our algorithm was deployed on a PC with an Intel® Duo 1.8GHz Core2™Duo 1.8 GHz processor and 2 GB RAM. The running time to process one neuron image is approximately 12 min. In Section 5.1 we explain the procedure used for validating the method, and show the results. Then, in Section 5.2 we explain the selection of parameters and quantitatively measure their effects on neuron reconstruction and spine identification.

5.1. Validation

The results of the spine identification procedure were validated on a data-set of 20 image stacks of pyramidal neuron cultures. Each 3D image contained multiple independent dendritic backbones with branching topology, an approximately fifteen spines per dendrite. Four expert users examined the original microscopy images and identified the dendritic spines. They then examined the dendrogram overlaid on the de-noised neuron 3D image (Fig. 11a and b), and the number of false positives (Type I error) and false negatives (Type II error) were tabulated across the users. The average sensitivity[6] and specificity[7] were then computed. These experiments were repeated for a number of parameter settings (see Fig. 15).

With regards to measuring the morphological features of dendrites, manual morphometry has been traditionally performed in 2D. It must be noted that there are no widely accepted methods of performing dendritic morphometry from 3D images. As a result, the manually measured values themselves are not accurate. Also, the measurements are subject to wide variation between the users. On our data-sets, the following morphological features were measured manually by the four users:

Length: Measured as the straight line distance between base of the spine and its apex, located in the 3D segmented volume,

Volume: Measured as the number of voxels in the spine enclosed by a plane marking the base of the spine in the 3D segmented volume,

Basal radius: Measured in the 2D maximum intensity projection of the dendrite at the base of the spine,

Maximum radius: Measured in the 2D maximum intensity projection of the dendrite at a point along the spine judged by the user to have maximum radius,

Minimum radius: Measured in the 2D maximum intensity projection of the dendrite at a point along the spine judged by the user to have minimum radius.

The dendritic spines were roughly grouped into two sets: short spines and long spines based on their length, and each set was evaluated separately. As the spine length reduces, the natural variability in locating the spine base and maximum (minimum) radius has a greater impact on the relative error (standard score). Therefore for shorter spines, the relative variability across human users in morphological measurements increases, making evaluation of the machine measurements less reliable. The accuracy of the machine morphometry for each feature (length, volume, basal radius, maximum radius, minimum radius) was computed using a MANOVA test (Krzanowski, 1988). Each spine in the set (short or long spines) was an independent variate, while the two samples to be tested were (a) the four manual measurements and (b) the four manual measurements along with the machine measurement. The difference between the means of the two samples was not found to be statistically significant (significance level of $\alpha = 0.05$). The Mahalanobis distance between the two sample means for each feature are tabulated in Table 1.

[3] MATLAB® is a product of The Mathworks Inc.

[4] Insight Segmentation and Registration Toolkit from the National Library of Medicine (NIH/NLM) (www.itk.org).

[5] Visualization Toolkit from Kitware Inc. (www.vtk.org).

[6] defined as $TP/(TP + FN)$, where TP: count of true positives, TN: count of true negatives, FP: count of false positives, FN: count of false negatives.

[7] defined as $TN/(TN + FP)$.

Table 1
MANOVA results for spine morphometry

Feature	Mahalanobis distance	
	Long spines	Short spines
Length	0.51	0.83
Volume	0.22	0.46
Basal radius	0.65	0.90
Maximum radius	0.42	0.77
Minimum radius	0.37	0.68

Table 2
List of parameters and their optimal values

Parameter	Explanation	Values	Type
λ_{nl}	Non-linear diffusion conductance (Eq. (2))	0.5–1.5	–
N_{nl}	Non-linear diffusion number of time-steps	15	–
σ	Regularization scale (Eq. (4))	2–5	Setup
Ω	Segmentation threshold (Section 3.3)	50–200	–
τ_{gac}	Level-set time-steps	5–10	Setup
ϑ	Floating fragment max volume ($\times 10^3$) (Section 3.4)	5–1000	–
α	Surface fairing smoothing (Eq. (9))	0.3–0.7	–
μ	Surface fairing expansion (Eq. (10))	−0.1 to −0.4	–
N	Surface fairing steps (Eq. (11))	20–80	–
ρ	Surface decimation factor	0.90–0.96	Data-set
θ	Skeletonisation strictness (Section 4.2)	0.0 to −0.2	Data-set
Γ_{min}	Minimum spine length (Section 4.3)	12 ± 3	Data-set
Γ_{max}	Maximum spine length (Section 4.4)	38 ± 5	Data-set
R_{max}	Maximum spine radius (Section 4.4)	17 ± 5	Data-set

Given the small size of each sample (4 or 5 data-points), the Mahalanobis distances observed are reasonable. The larger error in length measurement is explained by the fact that manually it was measured as the straight line distance between base to apex. Similarly, the larger errors in radius measurements are because the manual verification was done in 2D. In all cases, it was observed that the machine measurements were slightly larger than the manual measurements, confirming this conclusion.

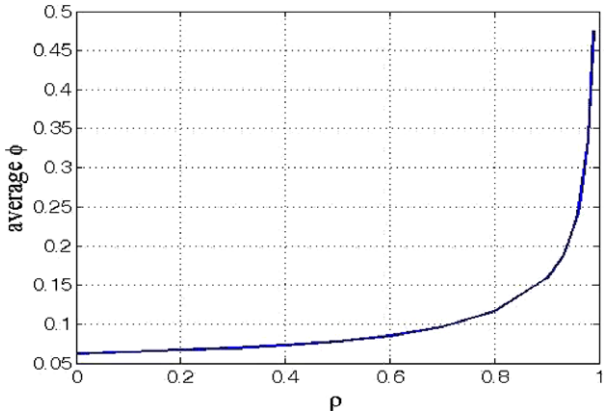

Fig. 14. Average dihedral angle $\overline{\phi}_{\mathbb{M}}$ of dendrite surface mesh $\mathbb{M}$ vs. decimation factor ρ.

5.2. Parameter selection

Table 2 gives a list of the parameters in the pipeline, their optimal values, and the conditions under which they need to be tuned. Parameters of type "*Setup*" depend upon either the properties of the image acquisition process or the digitization process and would need to be changed only if this setup were to change. The parameters of type "*data-set*" depend on the characteristics of

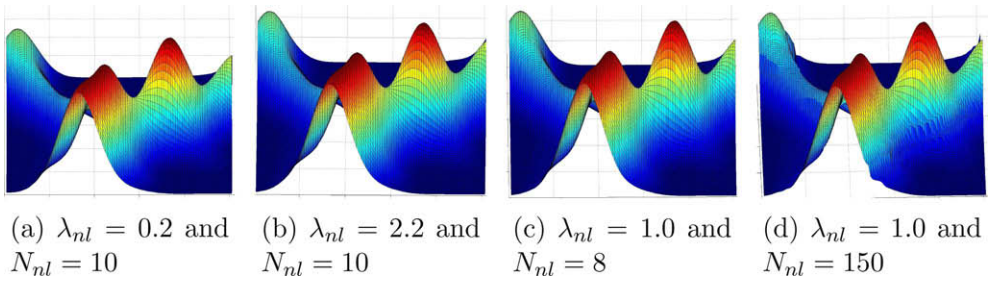

Fig. 12. The effect of conductance λ_{nl} and number of time-steps N_{nl} on non-linear diffusion filtering of the neuron images: (a)–(d) Monge map (intensity map) for a 2D section of the 3D volume for varying values of λ_{nl} and N_{nl}.

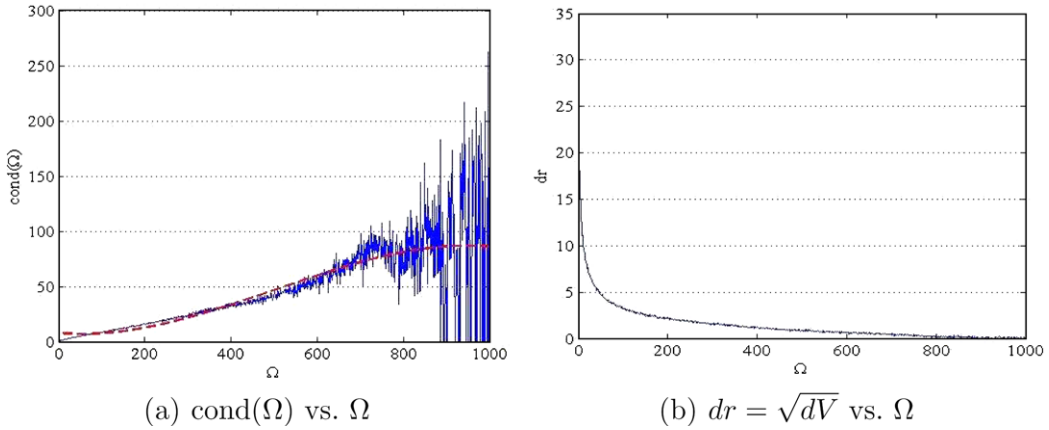

Fig. 13. (a) Condition number with respect to Ω. The red line is a least-squares fit of the data and (b) approximate radius (in voxels) of the segmented volume with respect to Ω.

the neurons being analyzed, and can be kept the same for the entire data-set.

Edges with gradient value less than the diffusion conductance parameter λ_{nl} are smoothed while those above it are enhanced. Given the nature of the intensity field in a neuron image, we find a large difference between the gradients of valid edges and noise edges, and the smoothing step is not sensitive to the selection of λ_{nl}. Perona and Malik (1990) have shown that edges remains stable over a long period of time in the non-linear diffusion process, while the solution gradually converges to a steady state. We too observe that after 10 iterations, most of the noise is removed and the edges remain stable up to 200 iterations (Fig. 12). Consequently, the number of time-steps N_{nl} is fixed at 15 for all data-sets. The regularization scale σ depends upon the physical spacing between the voxels. For our microscopy setup, a value of σ between 2 and 5 was found to be satisfactory.

The condition number cond(Ω) (Eq. (5)) quantifies the sensitivity of the segmentation threshold Ω (Section 3.3). As seen in Fig. 13a, the condition number is low for $\Omega < 300$, and once the background tissue is segmented out ($\Omega > 50$), the approximate radius (in voxels) of the segmented neuron stabilizes (Fig. 13b). As a result of this large range in acceptable values of Ω, the algorithm is robust with respect to it and it does not have to be fine tuned.

The τ_{gac} parameter enforces a proximity constraint (Section 3.4), between valid floating spine heads and the dendritic backbone, and depends upon the characteristics of the imaging process and the underlying tissue, which are responsible for this disconnect. We find that $\tau_{gac} \approx 5$ is appropriate for all cases in our data-set. The volumes of the smallest dendrite (>120,000 voxels) and the largest floating tissue fragment ($\ll$5000 voxels) differ by two orders of magnitude, and therefore the volume threshold ϑ has a lot of slack in its selection.

The two parameters α and μ of the surface fairing step affect the pass-band and stop-band of the surface frequencies (Section 4.1). The parameter N controls the sharpness of the cut-off frequency. Their value can be kept fixed and does not have to be tuned. This is because the scale of the noise in the iso-surface is many orders of magnitude smaller than the scale of the features (spines) in the neurons, as can be clearly seen from Fig. 9a.

To select the best quadric error decimation factor ρ we use the average dihedral angle $\overline{\phi}_{\mathbb{M}}$ of the surface mesh $\mathbb{M}$ to quantify smoothness. Fig. 14 shows the value of $\overline{\phi}_{\mathbb{M}}$ with respect to the decimation factor ρ. Initially, as number of polygons in the original mesh starts reducing, $\overline{\phi}_{\mathbb{M}}$ of the remains fairly constant. However, after a certain percentage of the faces are removed, the surface begins to develop sharp edges and $\overline{\phi}_{\mathbb{M}}$ begins to rise sharply. The optimal value of ρ is in the region, where the knee point occurs. We have found a factor of 0.8 to 0.95 to produce good results.

The skeletonisation strictness parameter θ (Section 4.2), the spine length thresholds $\Gamma_{\min}$ (Section 4.3), $\Gamma_{\max}$ and maximum spine radius threshold $R_{\max}$ (Section 4.4) work in conjunction to directly influence the sensitivity and specificity of the spine identification process. The values of $\Gamma_{\min}, \Gamma_{\max}$ and $R_{\max}$ serve to impose

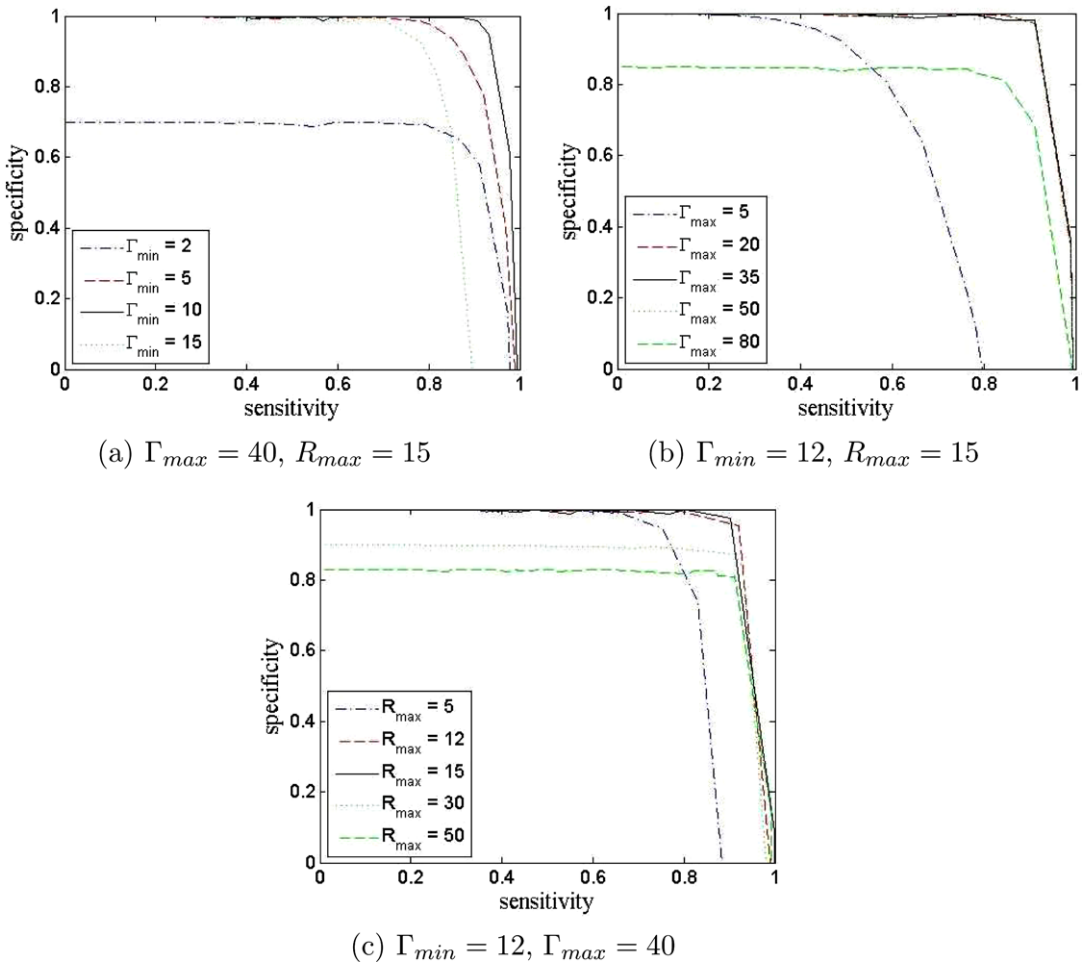

(a) $\Gamma_{max} = 40,\ R_{max} = 15$

(b) $\Gamma_{min} = 12,\ R_{max} = 15$

(c) $\Gamma_{min} = 12,\ \Gamma_{max} = 40$

Fig. 15. ROC curves with respect to θ: (a) ROC curve for different values of $\Gamma_{\min}$; (b) ROC curve for different values of $\Gamma_{\max}$; and (c) ROC curve for different values of $R_{\max}$.

biologically meaningful constraints on the size of the spines, and are known *a priori*, while θ is set constant for the entire data-set. From the ROC[8] curves (Fig. 15) we observe the effect of variations in $\Gamma_{\min}, \Gamma_{\max}$ and $R_{\max}$ on the accuracy of the algorithm. We obtained optimal sensitivity (0.953) and specificity (0.901) at $\theta = 0.05$, $\Gamma_{\min} = 12$, $\Gamma_{\max} = 38$ and $R_{\max} = 17$ for our data-set.

6. Conclusion

In this paper, we have presented a method to robustly reconstruct neuronal dendrites in 3D and to accurately identify spines on the dendrites. We developed a surface representation of the neuron that is compact and allows us control over the smoothness of the reconstructed surface. The curve-skeleton of the neuron was extracted using a procedure based on the medial geodesic function, which is robust to noise and correct in the sense of topological preservation. Based on the surface representation and the curve-skeleton we could accurately detect spines and measure spine length, volume, radius and other morphological features. Such features, in combination with other experimental information, may help researchers to delineate the mechanisms and pathways of neurological conditions such as Alzheimer's disease and tuberous sclerosis complex syndrome. The parameters of our method are easy to understand and biologically motivated. The accuracy of the system for spine identification and morphometry was demonstrated through a quantitative evaluation. We also presented a detailed study of the effect of the various parameters on the image processing, reconstruction and spine identification procedures that shows the robustness of this approach to parameter selection.

We are investigating methods to use the rich description of neuronal structure presented here to track dendrites and spines over time and study their morphological changes. By combining the 3D reconstruction algorithm with image registration, we plan to study the relationship between the changes in spines and their synaptic formations dynamically in order to uncover potentially new mechanisms of neuronal networks and functions. We also believe that the processing time for one neuron image can be reduced by a factor of $2\times$ by optimizing our algorithm and implementing it entirely in C/C++.

Acknowledgements

We thank Prof. Karen Zito, UC-Davis, for providing the neuron image data-sets. We also thank Prof. T.K. Dey, OSU, for inspiring this approach and for the use of their software. This work is partially supported by NIH-BISTI Grant (P20 EB000591-02). The work of X. Xu and S.T.C. Wong is supported in part by the Harvard NeuroDiscovery Center and Repair, Harvard Medical School, and Functional and Molecular Imaging Center, Department of Radiology, Brigham and Women's Hospital.

References

Al-Kofahi, K.A., Lasek, S., Szarowski, D.H., Pace, C.J., Nagy, G., Turner, J.N., Roysam, B., 2002. Rapid automated three-dimensional tracing of neurons from confocal image stacks. IEEE Transactions on Information Technology in Biomedicine 6 (2), 171–187.

Benshalom, G., White, E.L., 1988. Dendritic spines are susceptible to structural alterations induced by degeneration of their presynaptic afferents. Brain Research 443, 377–382.

Breu, Heinz, Gil, J., Kirkpatrick, D., Werman, M., 1997. Linear time euclidean distance transform algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence 17 (6), 529–533.

Can, A., Shen, H., Turner, J., Tanenbaum, H., Roysam, B., 1999. Rapid automated tracing and feature extraction from live high-resolution retinal fundus images using direct exploratory algorithms. IEEE Transactions on Information Technology in Biomedicine 3, 125–138.

Carlbom, I., Terzopoulos, D., Harris, K.M., 1994. Computer-assisted registration, segmentation, and 3d reconstruction from images of neuronal tissue sections. IEEE Transactions on Medical Imaging 13 (2), 351–361.

Catté, F., Lions, P.-L., Morel, J.-M., Coll, T., 1992. Image selective smoothing and edge detection by nonlinear diffusion. SIAM Journal of Numerical Analysis 29 (1), 182–193.

Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M., 1989. Detection of blood vessels in retinal images using two-dimensional matched filters. IEEE Transactions Medical Imaging 8, 263–269.

Choi, S.W., Seidel, H.-P., 2002. Linear one-sided stability of mat for weakly injective domain. Journal of Mathematical Imaging and Vision 17 (3), 237–247.

Clark, T.M., Freeman, W.R., Goldbaum, M.H., 1992. Digital overlay of fluorescein angiograms and fundus images for treatment of subretinal neovascularization. Journal of Retinal Vitreous Diseases 2 (12), 118–126.

Cohen, A.R., Roysam, B., Turner, J.N., 1994. Automated tracing and volume measurements of neurons from 3-D confocal fluorescence microscopy data. Journal of Microscopy 173 (2), 103–114.

Comery, T.A., Harrisdagger, J.B., Willems, P.J., Oostra, B., Irwin, S.A., Weiler, I.J., Greenough, W.T., 1997. Abnormal dendritic spines in fragile X knockout mice: maturation and pruning deficits. PNAS 94, 5401–5404.

Cornea, N., Silver, D., Min, P., 2005. Curveskeleton applications. In: IEEE Conference on Visualization, pp. 95–102.

Dey, T.K., Sun, J., 2006a. Curveskel software. <http://www.cse.ohio-state.edu/>.

Dey, T.K., Sun, J., 2006b. Defining and computing curve-skeletons with medial geodesic function. In: Eurographics Symposium on Geometry Processing, pp. 143–152.

Dima, A., 2002. Computer Aided Image Segmentation and Graph Construction of Nerve Cells from 3d Confocal Microscopy Scans. Ph.D. Thesis, Technical University Berlin, Berlin, Germany.

Dima, A., Scholz, M., Obermayer, K., 2001. Automatic generation of 3d-graph representations of nerve cells from confocal microscope scans. Göttingen Neurobiology Conference, vol. 2. German Neuroscience Society, p. 1041.

Engert, F., Bonhoeffer, T., 1999. Dendritic spine changes associated with hippocampal long-term synaptic plasticity. Nature 399, 66–70.

Garland, M., Heckbert, P.S., 1998. Simplifying surfaces with color and texture using quadric error metrics. In: IEEE Conference on Visualization.

Garvey, C.F., Young, J., Simon, W., Coleman, P.D., 1973. Automated three-dimensional dendrite tracking system. Electroencephalography and Clinical Neurophysiology 35, 199–204.

Gerig, G., Koller, T., Székely, G., Brechbűhler, C., Kűbler, O., 1993. Symbolic description of 3-d structures applied to cerebral vessel tree obtained from MR angiography volume data. Processing in Medical Imaging (687), 94–111.

Glantz, L.A., Lewis, D.A., 2000. Decreased dendritic spine density on prefrontal cortical pyramidal neurons in schizophrenia. Archives of General Psychiatry 57 (1), 65–73.

Glaser, J.R., Glaser, E., 1990. Neuron imaging with neurolucida – a PC-based system for image combining microscopy. Computerized Medical Imaging and Graphics 14, 307–317.

Goldbaum, M.H., Katz, N., Chaudhuri, S., Nelson, M., Kube, P., 1990a. Digital image processing for ocular fundus images. Ophthalmology Clinics North America 3 (3), 447–466.

Goldbaum, M.H.V., Kouznetsova, B.L., Cote, W.E.H., Nelson, M., 1990b. Digital image processing for ocular fundus images. Ophthalmology Clinics North America 3 (3), 447–466.

Harris, K.M., Stevens, J., 1988. Study of dendritic spines by serial electron microscopy and three-dimensional reconstructions. Neurology and Neurobiology 37, 179–199.

Harris, K.M., Jensen, F.E., Tsao, B., 1992. Three-dimensional structure of dendritic spines and synapses in rat hippocampus (CA1) at postnatal day 15 and adult ages: implications for the maturation of synaptic physiology and long-term potentiation. Journal of Neuroscience 12, 2685–2705.

He, W., Hamilton, T.A., Cohen, A.R., Holmes, T.J., Turner, J.N., Roysam, B., 2003. Automated three-dimensional tracing of HRP stained neurons from a stack of brightfield optical slices. Microscopy and Microanalysis 9, 296–310.

Hering, H., Sheng, M., 2001. Dendritic spine: structure, dynamics and regulation. Neuron 2, 880–888.

Herzog, A., Krell, G., Michaelis, B., Wang, J., Zuschratter, W., Braun, K., 1997. Restoration of three-dimensional quasi-binary images from confocal microscopy and its application to dendritic trees. In: SPIE Three-Dimensional Microscopy: Image Acquisition and Processing IV, pp. 146–157.

Herzog, A., Krell, G., Michaelis, B., Zuschratter, W., 1998. Tracking on tree-like structures in 3-d confocal images. 3-D and Multidimensional Microscopy: Image Acquisition and Processing 3261, 165–176.

Kasai, H., Matsuzaki, M., Noguchi, J., Yasumatsu, N., Nakahara, H., 2003. Structure–stability–function relationships of dendritic spines. Trends Neuroscience 26 (3), 360–368.

Koh, I.Y.Y., 2001. Automated Recognition Algorithms for Neural Studies. Ph.D. Thesis, Stony Brook University.

Koh, I.Y.Y., Lindquist, W.B., Zito, K., Nimchinsky, E.A., Svoboda, K., 2002. An image analysis algorithm for dendritic spines. Neural Computation 14 (6), 1283–1310.

Krzanowski, W., 1988. Principles of Multivariate Analysis. Oxford University Press.

Ma, W., Wu, F., Ouhyoung, M., 2003. Skeleton extraction of 3d objects with radial basis functions. IEEE SMA.

[8] Receiver operating characteristic.

Malladi, R., Sethian, J.A., Vemuri, B.C., 1995. Shape modeling with front propagation: a level set approach. IEEE Transactions on Pattern Analysis and Machine Intelligence 17 (2), 158–175.

Mosaliganti, K., Janoos, F., Xu, X., Machiraju, R., Huang, K., Wong, S., 2006. Temporal matching of dendritic spines in confocal microscopy images of neuronal tissue sections. In: MICCAI Workshop on Medical Image Analysis with Applications in Biology, pp. 106–113.

Moser, M.B., Trommald, M., Anderson, P., 1994. An increase in dendritic spine density on hippocampal CA1 pyramidal cells following spatial learning in adult rats suggests the formation of new synapses. National Academy of Science 91, 12673–12675.

Nielson, G.M., Hamann, B., 1991. The asymptotic decider: resolving the ambiguity in marching cubes. In: IEEE Conference on Visualization.

Nimchinsky, E.A., Sabatini, B.L., Svoboda, K., 2002. Structure and function of dendritic spines. Annual Review in Physiology (64), 313–352.

Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence 12 (7), 629–639.

Rusakov, D.A., Stewart, M.G., 1995. Quantification of dendritic spine populations using image analysis and a tilting disector. Journal of Neuroscience Methods 60, 11–21.

Schmitt, S., Dima, A., Scholz, M., Obermayer, K., 2001. Automatic three-dimensional reconstruction of neurons from confocal images. In: Göttingen Neurobiology Conference.

Schmitt, S., Evers, J.-F., Scholz, C.D.M., Obermayer, K., 2004. New methods for the computer-assisted 3d reconstruction of neurons from confocal image stacks. NeuroImage.

Spires, T.L., Meyer-Luehmann, M., Stern, E.A., McLean, P.J., Skoch, J., Nguyen, P.T., Bacskai, B.J., Hyman, B.T., 2005. Dendritic spine abnormalities in amyloid precursor protein transgenic mice demonstrated by gene transfer and intravital multiphoton microscopy. Journal of Neuroscience 25 (31), 7278–7287.

Stevens, J.K., Trogadis, J., 1984. Computer-assisted reconstruction from serial electron micrographs: a tool for the systematic study of neuronal form and function. Advances in Cellular Neurobiology 5, 341–369.

Taubin, G., 1995. A signal processing approach to fair surface design. In: SIGGRAPH Conference on Computer Graphics and Interactive Techniques, pp. 351–358.

Tavazoie, S.F., Alvarez, V.A., Ridenour, D.A., Kwiatkowski, D.J., Sabatini, B.L., 2005. Regulation of neuronal morphology and function by the tumor suppressors Tsc1 and Tsc2. Nature Neuroscience 8, 1727–1734.

Tyrrell, J., Roysam, B., di Tomaso, E., Tong, R., Brown, E., Jain, R., 2006. Robust 3-d modeling of tumor microvasculature using superellipsoid. In: IEEE International Symposium on Biomedical Imaging: Macro to Nano, vol. 6, pp. 185–188.

Uehara, C., Colbert, C., Saggau, P., Kakadiaris, I.A., 2004. Towards automatic reconstruction of dendrite morphology from live neurons. In: IEEE Conference of the Engineering in Medicine and Biology Society.

Urban, E.S., O'Malley, S.M., Walsh, B., Santamara-Pang, A., Saggau, P., Colbert, C., Kakadiaris, I.A., 2006. Automatic reconstruction of dendrite morphology from optical section stacks. In: Workshop on Computer Vision Approaches to Medical Image Analysis, pp. 190–201.

Watzel, R., Braun, K., Hess, A., Scheich, H., Zuschratter, W., 1995. Detection of dendritic spines in three-dimensional images. In: Deutsche Arbeitsgemeinschaft für Mustererkennung Symposium, pp. 160–167.

Weaver, C.M., Hof, P.R., Wearne, S.L., Lindquist, W.B., 2004. Automated algorithms for multiscale morphometry of neuronal dendrites. Neural Computation 16 (7), 1353–1383.

Weickert, J., 1996. Anisotropic Diffusion in Image Processing. Ph.D. Thesis, Universität Kaiserslautern.

Wilson, C.J., Murakami, F., Katsumaru, H., Tsukahara, N., 1987. Dendritic and somatic appendages of identified rubrospinal neurons of the cat. Neuroscience 22, 113–130.

Witkin, A., 1983. Scale-space filtering. In: International Joint Conference on Artificial Intelligence. Karlsiuhe, West Germany, pp. 1019–1021.

Yuste, R., Bonhoeffer, T., 2001. Morphological changes in dendritic spines associated with long-term synaptic plasticity. Annual Reviews in Neuroscience 24, 1071–1089.

Zito, K., Knott, G., Shepherd, G., Shenolikar, S., Svoboda, K., 2004. Induction of spine growth and synapse formation by regulation of the spine actin cytoskeleton. Neuron 44 (2), 321–334.