CrossMark

Stability, structure and scale: improvements in multi-modal vessel extraction for SEEG trajectory planning

Maria A. Zuluaga[1] · Roman Rodionov[2,3] · Mark Nowell[2,3] · Sufyan Achhala[2] ·
Gergely Zombori[1] · Alex F. Mendelson[1] · M. Jorge Cardoso[1] ·
Anna Miserocchi[2,3] · Andrew W. McEvoy[2,3] · John S. Duncan[2,3] ·
Sébastien Ourselin[1]

Abstract

Purpose Brain vessels are among the most critical landmarks that need to be assessed for mitigating surgical risks in stereo-electroencephalography (SEEG) implantation. Intracranial haemorrhage is the most common complication associated with implantation, carrying significantly associated morbidity. SEEG planning is done pre-operatively to identify avascular trajectories for the electrodes. In current practice, neurosurgeons have no assistance in the planning of electrode trajectories. There is great interest in developing computer-assisted planning systems that can optimise the safety profile of electrode trajectories, maximising the distance to critical structures. This paper presents a method that integrates the concepts of scale, neighbourhood structure and feature stability with the aim of improving robustness and accuracy of vessel extraction within a SEEG planning system.

Methods The developed method accounts for scale and vicinity of a voxel by formulating the problem within a multi-scale tensor voting framework. Feature stability is achieved through a similarity measure that evaluates the multi-modal consistency in vesselness responses. The proposed measurement allows the combination of multiple images modalities into a single image that is used within the planning system to visualise critical vessels.

Results Twelve paired data sets from two image modalities available within the planning system were used for evaluation. The mean Dice similarity coefficient was 0.89 ± 0.04, representing a statistically significantly improvement when compared to a semi-automated single human rater, single-modality segmentation protocol used in clinical practice (0.80 ± 0.03).

Conclusions Multi-modal vessel extraction is superior to semi-automated single-modality segmentation, indicating the possibility of safer SEEG planning, with reduced patient morbidity.

Keywords Computer-assisted planning system · Vessel extraction · Depth electrode insertion · Image-guided surgery · Multi-modal segmentation

✉ Maria A. Zuluaga
M.ZuluagaValencia@cs.ucl.ac.uk; maria.zuluaga@ucl.ac.uk

[1] Translational Imaging Group, CMIC, University College London, London, UK

[2] Department of Clinical and Experimental Epilepsy, UCL IoN, London, UK

[3] National Hospital for Neurology and Neurosurgery (NHNN), London, UK

Introduction

The primary goal of epilepsy surgery is to remove the epileptogenic zone, the minimum amount of cortex that must be resected to produce seizure freedom [3,16]. As the epileptogenic zone may not be associated with a clear structural abnormality, intracranial electrodes must be used to record the area of the brain where seizures start, known as the seizure-onset zone (SOZ) [3]. Stereo-electroencephalogra-

phy (SEEG) is the recording of the brain electrical activity by depth electrodes implanted into the brain parenchyma to precisely identify the SOZ. The major complication associated with SEEG implantation is intracranial haemorrhage with a risk that ranges from 0.6 to 2.7 % [13] and reported morbidity and mortality of 5.6 % [2] and 1 % (or less) [13], respectively. To reduce the risk of this and other associated complications, it is necessary to identify electrode trajectories with adequate cortical coverage that pass through safe avascular planes. This is done through pre-operative SEEG planning.

In recent years, there has been great interest in the development of computer-assisted planning systems for optimising intracranial depth electrode insertion [1,4,5,18]. These methods require the effective extraction of critical brain landmarks (such as vessels, ventricles and sulci), with high accuracy and robustness. Despite this requirement, the techniques used to extract the landmarks are very general and may not be the best for this application. While some efforts have been made to improve them [1,4], greater improvements yet may be obtained using domain-specific knowledge. Furthermore, as pointed by Du et al. [4], the evaluation or validation of the extraction of these brain landmarks is not included in these studies, which hinders the identification of potential problems of the currently used techniques. In this work, we specifically address the extraction of the intracranial vasculature within an SEEG planning system.

A key part of the landmark identification is the vessel extraction. This, despite years of research [9], remains a challenging problem. For this reason, Essert et al. [5], avoid dealing with it directly. Instead, they segment the cortical sulci under the hypothesis that vessels of interest are generally located there. Existing methods of vessel extraction still tend to suffer from discontinuities (caused by low intensity due to partial volume effects and noise) and false-positives; both of these are undesirable as they can lead to invalid or suboptimal trajectories. A common solution, adopted by Bériault et al. [1] and Du et al. [4], is the use of a vesselness filter [6,10,17] that enhances voxels within tubular structures. These filters have been very successful thanks to the inclusion of multiple spatial scales within their formulation, but lack information about the surrounding structures. Also, despite increased access to multi-modal images, particularly within computer-assisted planning systems, very few methods have exploited the information complementarity to improve vessel extraction. Passat et al. [15] combined multiple MR sequences to segment the superior sagittal sinus, but their vessel extraction was only performed on a single image, with a second modality used to provide a priori anatomical information of the brain. More recently, Bériault et al. [1] used susceptibility weighted imaging (SWI) and time-of-flight (TOF) images within a planning system to segment veins and arteries, but they do so in a separate fashion and leave

it up to the human planner to combine the information from the two modalities. Elsewhere, Hu et al. [8] have proposed a bi-modal approach to 2D retinal vessel extraction using a k-NN classifier trained on features from both modalities.

In this paper, we present a novel method that integrates the concepts of scale, neighbourhood structure and feature stability with the aim of improving the robustness and accuracy of vessel extraction within a computer-assisted SEEG planning system [19]. The method accounts for both the scale and vicinity of a voxel by formulating the problem within a multi-scale tensor voting framework. Feature stability is achieved by introducing a similarity measure that accounts for the multi-modal consistency in the vesselness responses. The proposed measurement enables the combination of multiple responses into a single image that is used within the planning system to visualise critical vessels. This article is an extended version of Zuluaga et al. [20]; it explains the methodology in more detail, extends the original method to cope with an arbitrary number of images and proposes a new way to consider multiple scales within the tensor voting framework to improve the computational speed. A new set of validation experiments is also included.

Method

The tensor voting framework [11] is a robust technique for extracting structures from a cloud of points. It is based on the principle that a set of unconnected tokens (i.e. points) can exchange information within a neighbourhood that allows one to infer the geometric structure in which a token lies. In 3D, it provides a way to estimate the likelihood that a token lies on a surface, curve or junction (as opposed to being just noise).

Tensor voting consists of three stages: token initialisation, where token locations are identified and their tensors assigned; tensor voting, where tensors from a token and its neighbours are combined; and analysis of voting results. In order to give our method feature stability and scale invariance, we add a data fusion step and then embed the different components into a multi-scale framework. A diagram illustrating the proposed method is shown in Fig. 1.

Multi-scale framework and voxel selection

In previous work [20], we formulated the tensor voting approach into a multi-scale framework by evaluating the response of different scales at the data fusion stage and retaining the maximum response through scales. This consists in performing steps 2–5 in Fig. 1 for every scale and image modality. As the tensor vote has to be computed every time, this can be very computationally intensive. To reduce the computational burden of the method, we now perform the
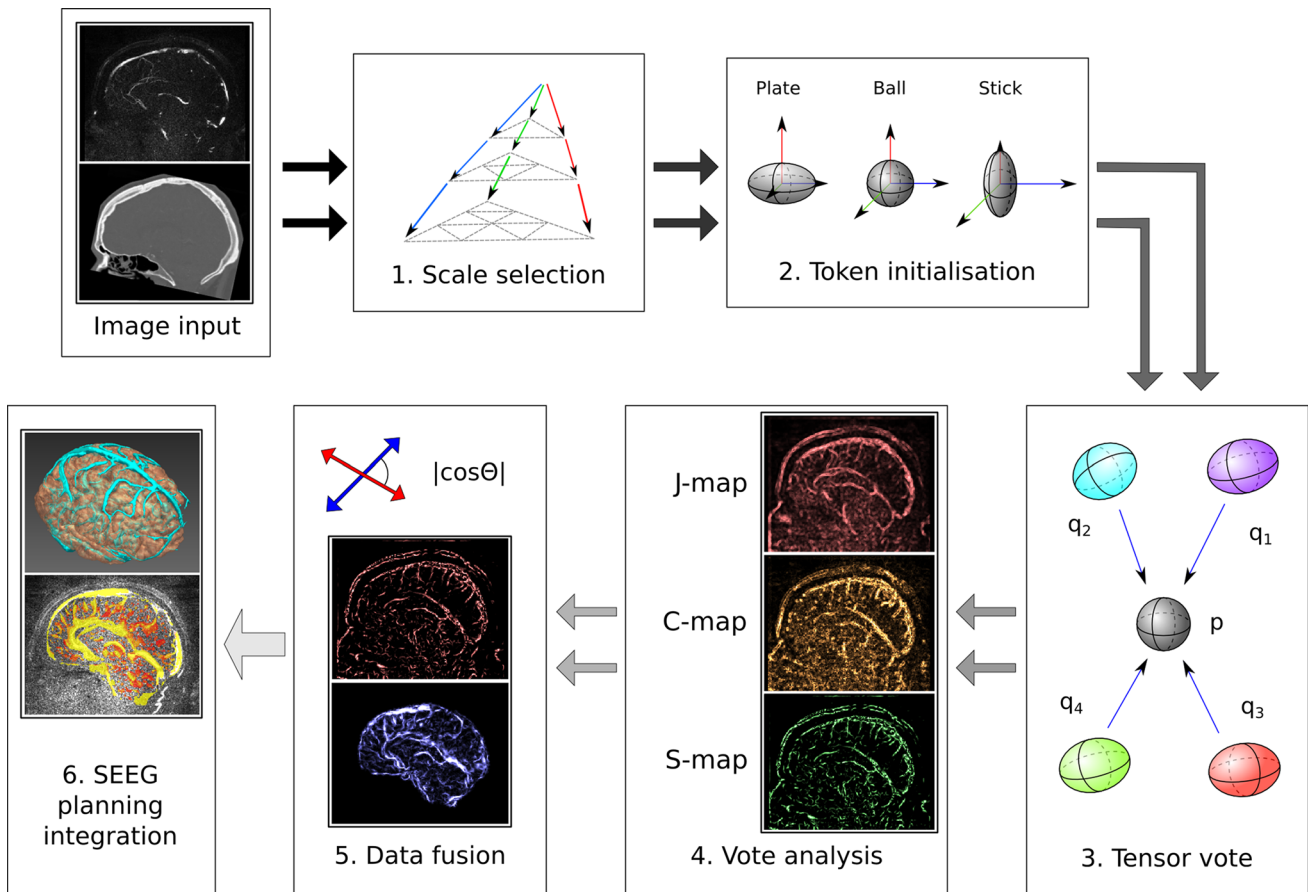
Fig. 1 Vessel extraction *diagram*. After optimal scale selection, *images* are converted into tokens through analysis of the Hessian matrix. After voting, the resulting saliency *maps* are combined using the cosine between the vectors defining orientation. The resulting probability *map* is then visualised in the planning system

multi-scale analysis before the tensor voting step, rather than at the final stage.

Let $I(\mathbf{x})$ be an image and $H_\sigma(\mathbf{x})$ be the Hessian matrix at voxel $\mathbf{x}$ at scale σ. By means of the eigenvalues κ_i and eigenvectors $\mathbf{v_i}$ of $H_\sigma(\mathbf{x})$, a vesselness function $\nu(\mathbf{x}, \sigma)$ [10] is computed at a range of scales $[\sigma_{\min}, \sigma_{\max}]$. An "optimal scale" image $M(\mathbf{x})$ is constructed by finding the value of σ with the maximum vesselness response at each voxel:

$$M(\mathbf{x}) = \sigma^*, \quad \sigma^* = \underset{\sigma_{\min} \leq \sigma \leq \sigma_{max}}{\arg \max} \ \nu(\mathbf{x}, \sigma). \quad (1)$$

While the Hessian and its eigensystem are computed across the full range of scales, all further calculations involving scale analysis are performed using the scale defined by $M(\mathbf{x})$. For simplicity, we denote the Hessian at the locally optimal scale as $H(\mathbf{x})$.

We derive from this eigensystem the information about which voxels are likely to belong to a vessel structure. The eigenvalues of $H(\mathbf{x})$ are ordered by their absolute value, so that $|\kappa_3| \geq |\kappa_2| \geq |\kappa_1|$. For bright vessels on a dark back-

ground, a voxel is expected to belong to a vessel only if $\kappa_2 < 0$ and $\kappa_3 < 0$ [10]. Based on this information, voxels that do not satisfy this condition for at least one scale can be discounted in the further analysis. Voxels satisfying the condition go on to take part in the tensor voting, with their Hessian computed at the locally optimal scale. Under the voting formalism, a voxel $\mathbf{x}$ considered for tensor voting is called a token and is denoted as $\mathbf{p}$.

Token initialisation

The information contained in a token $\mathbf{p}$ is encoded on a 3D second-order, symmetric, nonnegative definite tensor $\mathbf{T}$. According to the spectrum theorem, $\mathbf{T}$ can be expressed as the linear combination of three tensors:

$$\mathbf{T} = (\lambda_1 - \lambda_2)(\mathbf{e_1}\mathbf{e_1}^T) + (\lambda_2 - \lambda_3) \sum_{i=1}^{2} \mathbf{e_i}\mathbf{e_i}^T + \lambda_3 \sum_{i=1}^{3} \mathbf{e_i}\mathbf{e_i}^T, \quad (2)$$

Table 1 Eigenvalues λ_i and eigenvectors $\mathbf{e_i}$ of $\mathbf{T}$ for each initial orientation initialisation

Initialisation	Eigenvalues λ_i	Eigenvectors $\mathbf{e_i}$
No preferred orientation	1: $\lambda_1 = \lambda_2 = \lambda_3 = K$	$\mathbf{e_i} = (e_{ij}),\, e_{ij} = \delta_{ij}$
	2: $\lambda_1 = \lambda_2 = \lambda_3 = I(\mathbf{p})$	
	3: $\lambda_1 = \lambda_2 = \lambda_3 = \nu(\mathbf{p})$	
Hessian matrix	1: $[\lambda_1, \lambda_2, \lambda_3] = [\lvert\kappa_1\rvert^{-1}, \lvert\kappa_2\rvert^{-1}, \lvert\kappa_3\rvert^{-1}]$	$[\mathbf{e_1}, \mathbf{e_2}, \mathbf{e_3}] = [\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}]$
	2: $[\lambda_1, \lambda_2, \lambda_3] = [\nu(\mathbf{p}), 0, 0]$	
Structure tensor	$(\lambda_1 - \lambda_2) = \lVert \bigtriangledown I \rVert^2,\, \lambda_3 = 0$	$\mathbf{e_1}\mathbf{e_1}^T = \bigtriangledown I \bigtriangledown I^T,\, \mathbf{e_2} = \mathbf{e_3} = \mathbf{0}$

Note that when no preferred orientation is used, three different saliency measurements can be used. Similarly, two different configurations of the initial eigenvalues are proposed when using Hessian matrix analysis

where λ_i are the eigenvalues obtained from the eigendecomposition of $\mathbf{T}$, with $\lambda_1 > \lambda_2 > \lambda_3 \geq 0$, and $\mathbf{e_i}$ are the corresponding eigenvectors.

While the tensor is most commonly expressed as a 3×3 matrix, it can be viewed as a 3D ellipsoid whose shape describes the contribution of its different components in Eq. 2. These components are known as the stick tensor S, the plate tensor P and the ball tensor B. More precisely, through expansion of the second and third terms, they are defined as:

$$S = (\lambda_1 - \lambda_2)\left(\mathbf{e_1}\mathbf{e_1}^T\right)$$
$$P = (\lambda_2 - \lambda_3)\left(\mathbf{e_1}\mathbf{e_1}^T + \mathbf{e_2}\mathbf{e_2}^T\right)$$
$$B = \lambda_3\left(\mathbf{e_1}\mathbf{e_1}^T + \mathbf{e_2}\mathbf{e_2}^T + \mathbf{e_3}\mathbf{e_3}^T\right). \tag{3}$$

Each tensor component corresponds to a different type of structural information: the stick tensor represents an elongated ellipsoid encoding eccentricity with orientation $\mathbf{e_1}$, the plate tensor represents a disc-shaped structure with normal $\mathbf{e_3}$, and the ball tensor represents a round structure in which all orientations are equally probable. The scalar values associated with each tensor are the saliency measurements of "surfaceness" $(\lambda_1 - \lambda_2)$, "curveness" $(\lambda_2 - \lambda_3)$ and "junctionness" λ_3. Points with very small eigenvalues are regarded as noise.

The scalar information contained in a greyscale image needs to be encoded into a tensor that satisfies Eqs. 2–3 before it can be used within the tensor voting framework. Depending on which tensor is assigned to the tokens in the token intitialisation step, different aspects of image will inform the final vessel segmentation. In this paper, we explore the three alternatives presented below.

No preferred orientation

A common approach is to assign an isotropic, ball-shaped tensor to each of the tokens [11,12]. All three eigenvalues of a token's tensor have the same value, K. While the tensors at each token have the same shape, they can be given different magnitudes (values of K) according to their tokens' vessel salience score. In this way, tokens that are most likely to

be on the vessels are given greater importance in the final segmentation. In this work, we experiment with using both a token's vesselness, $\nu(\mathbf{p})$, or its image intensity, $I(\mathbf{p})$, as salience score/weight.

Hessian matrix-based orientation

The analysis of the eigensystem of the Hessian matrix provides information about the orientations of structures within an image [6]. Let us recall $\mathbf{v_i}$, the eigenvectors of $H(\mathbf{p})$ at the optimal scale, and $\lvert\kappa_3\rvert \geq \lvert\kappa_2\rvert \geq \lvert\kappa_1\rvert$, its eigenvectors. Tokens are initialised with a stick tensor S that has orientation $\mathbf{v_1}$ (the eigenvector associated with κ_1), the direction along the vessel. The saliency of the stick tensor is obtained by assigning the κ_i^{-1} to the corresponding λ_i. As with the ball tensor, it is possible to initialise the stick tensors with different weights reflecting an initial estimate of a token's vesselness. Section "Experiments and results" compares the results obtained when varying these values.

Structure tensor-based initialisation. Given $\bigtriangledown I(\mathbf{p})$,[1] the image gradient at a given token position $\mathbf{p}$, Moreno et al. [12] showed that the tensorised gradient, $\bigtriangledown I \bigtriangledown I^T$, can be used to initialise a stick tensor S at every token, giving a saliency measurement $\lVert \bigtriangledown I \rVert^2$.

Table 1 summarises the values that are assigned to the tensor's eigenvalues, λ_i, and eigenvectors, $\mathbf{e_i}$, depending on the type of used initialisation.

Tensor voting

After each token $\mathbf{p}$ is encoded as a tensor $\mathbf{T}$ in the form of Eq. 2, it propagates structural information to its neighbours in the form of a vote. Votes are combined through addition at every token to infer the type of structure going through it. More formally, the tensor voting at $\mathbf{p}$ is given by [12]:

---

[1] For the sake of simplicity, we will drop the index term and the image gradient at a given token position $\mathbf{p}$ will be expressed as $\bigtriangledown I$ for the remaining of this article.

Table 2 Feature maps 2-tuple $(s, \mathbf{w})$ definition

Map	s	$\mathbf{w}$
S-Map	$\lambda_1 - \lambda_2$	$\mathbf{e_1}$
C-Map	$\lambda_2 - \lambda_3$	$\mathbf{e_3}$
J-Map	λ_3	Arbitrary

$$TV(\mathbf{p}) = \sum_{\mathbf{q} \in \chi} \left(SV\left(\mathbf{v}, S_{\mathbf{q}}\right) + PV\left(\mathbf{v}, P_{\mathbf{q}}\right) + BV\left(\mathbf{v}, B_{\mathbf{q}}\right) \right), \tag{4}$$

where χ denotes the neighbourhood of $\mathbf{p}$, $\mathbf{q}$ a point belonging to χ, SV, PV and BV the stick, plate and ball votes cast to $\mathbf{p}$ by each component $S_{\mathbf{q}}$, $P_{\mathbf{q}}$, $B_{\mathbf{q}}$ of $\mathbf{q}$ and $\mathbf{v} = \mathbf{p} - \mathbf{q}$. The strength of the vote will be dependent on the norm of $\mathbf{v}$, as the influence of a point $\mathbf{q}$ should decay as its distance from $\mathbf{p}$ increases. Here, $q \in \chi$ is defined as the window of size $N_{\mathbf{p}}$ at every token with

$$N_{\mathbf{p}} = 2M(\mathbf{p}), \tag{5}$$

where $M(\mathbf{p})$ is the "optimal scale" image defined in Eq. 1.

The derivation of SV, PV and BV in Eq. 4 is extensive and beyond the scope of this paper. We therefore only present their formulation (see Appendix 6), mentioning that the tensor voting procedure can be regarded as a tensor convolution with a voting kernel which itself produces a tensor. The interested reader is referred to [11,12] for further details.

Voting analysis

As the result of the tensor voting is another tensor, it can be decomposed as in Eq. 2. From this decomposition, three feature vector maps, the surface (S-Map), the curve map (C-Map) and the junction map (J-Map), are constructed. A voxel of these maps contains a 2-tuple $(s, \mathbf{w})$, where s is a scalar indicating strength/saliency and $\mathbf{w}$ is a unit vector indicating direction. Table 2 summarises the values for $(s, \mathbf{w})$ within the different feature maps.

In the context of our problem, we are interested in the information provided by the S-map (first term of Eq. 2). For a given tuple, we interpret s as a consensus measurement of vesselness between a voxel and its neighbours and $\mathbf{w}$ as the direction along the vessel.

Data fusion

Given two D-dimensional vectors, the cosine of the angle between them is an index on the extent to which they are aligned. As vessels are well-oriented structures, the cosine of the direction vectors is a surrogate of vesselness consistency between different images. Given two sets of tuples $(s_1, \mathbf{w_1})$

and $(s_2, \mathbf{w_2})$ from vesselness maps obtained from two different modalities after voting analysis, with $\|\mathbf{w_1}\| = \|\mathbf{w_2}\| = 1$, it is possible to fuse the maps into a single one through the following expression:

$$\varphi(\mathbf{p}) = 0.5|\mathbf{w_1} \cdot \mathbf{w_2}|(s_1 + s_2), \tag{6}$$

where $\cdot$ denotes the dot product between the two vectors. By refactoring Eq. 6, the fusion can be extended to L different image modalities through:

$$\varphi(\mathbf{p}) = \frac{1}{L} \sum_{i=1}^{L} s_i \sum_{j=1}^{L} (1 - \delta_{ij})|\mathbf{w_i} \cdot \mathbf{w_j}|. \tag{7}$$

The fusion scheme is a measure that rewards consensus and punishes discord; the greater the angular distance between the different directions, the more the absolute value of the resulting output is reduced. When there is complete agreement between modalities, the output is simply the average. When their directions are perpendicular, the output drops to zero to reflect a complete lack of certainty.

SEEG planning system integration

Similar to Bériault et al. [1] and Du et al. [4], the resulting vessel probability map, $\varphi(\mathbf{p})$, is used as input of our computer-assisted planning system [19]. As the electrode-implanting trajectory needs to be further than a safety margin from the critical tissue (vessels in this case), the probability map serves as a measure of risk of crossing a vessel. Within the planning system, the probability map is converted into a 3D surface mesh object, coloured using a pre-defined landmark colour scheme [19] and displayed within the neuronavigation planning system along with other brain structures (Fig. 3).

Experiments and results

Data

Blood vessel images for the computer-assisted planning system [19] were acquired using 3D phase contrast MRI (3DPC) and CT angiograms (CTAs). For our experiments, we used twelve paired data sets of 3DPC and CTA available within the planning system (informed consent obtained from all the patients). The 3DPC data were acquired on a 1.5-T Siemens Avanto MR scanner with voxel size resolution $0.8593 \times 0.8593 \times 1 \, mm^3$ and velocity encoding of 5 cm/s in each direction. CTA images were acquired on a Siemens SOMATOM Definition AS+ scanner with voxel size resolution $0.4296 \times 0.4296 \times 0.75 \, mm^3$. During image acquisition,

the patient's head was immobilised by placing a pad between the head and the coil.

Gold standard generation

Three different observers (a neurosurgical trainee, a physicist with 8-year experience in clinical neuroimaging and a master student trained for the task) segmented blood vessels structures following the protocol typically used in clinical practice for SEEG planning in the absence of a computer-assisted planning system. The semi-automated annotation procedure followed the steps described hereafter:

1. An intracranial space mask was applied to the CTA image to remove skull and to the 3DPC image to remove extracranial blood vessels.
2. Masked CTA and 3DPC were separately thresholded to give an initial estimate of the vessels. The threshold was defined by visually evaluating the resulting segmentation and determining whether noise and blood vessels were easy to distinguish and differentiate from each other with minimal manual cleaning.
3. Small isolated clusters were removed based on diameter size within MeshLab. The observers varied the threshold until they considered the segmented result satisfactory through visual inspection. Afterwards, large noise (e.g. calcifications) was removed by manually editing the images using MeshLab.

The six segmentations of the observers were combined into a consensus agreement through a voting strategy in a similar fashion as Hameeteman et al. [7].

Validation scheme

The proposed algorithm was evaluated on the twelve affinely co-registered [14] data sets using ten different scales between $\sigma_{\min} = 1.0$ and $\sigma_{\max} = 4.5$ distributed equally in a log space.

For a quantitative evaluation, each of the twelve vessel images S was compared to the consensus agreement M using the Dice similarity coefficient (DSC):

$$\mathrm{DSC} = \frac{2 \times |S \cap M|}{|S + M|}, \tag{8}$$

where the intersection operation is the voxel-wise minimum operation and $|\cdot|$ is the integration of the voxel values over the complete image [7].

We used the DSC to assess the performance of our method and that one of each observer w.r.t the consensus when doing a semi-automated segmentation with a single modality as is done in clinical routine. We could then compare the accuracy of the proposed method to current practice.

Evaluation of different initialisation strategies

In Section "Token initialisation", we presented different strategies to initialise the token's tensors. Initialisation involves the definition of an initial orientation and the saliency of the tensors. In order to determine the effects of initialisation in vessel extraction, we evaluated the performance of the six different initialisation configurations (Table 1).

No preferred orientation

Three different saliency measurements were estimated to initialise ball tensors: a constant value for every tensor $K = 1$, the image intensity at each token, $\lambda_{1-3} = I(\mathbf{p})$, and a vesselness measurement at each token $\lambda_{1-3} = \nu(\mathbf{p})$. We chose the vesselness measurement proposed by Manniesing et al. [10] due to its smoothness properties. In order to compute the vesselness function, we followed the guidelines reported in the original publication [10].

Reported DSCs were 0.42 ± 0.04, 0.79 ± 0.03 and 0.89 ± 0.04 for constant value, intensity-based and vesselness measurement initialisation, respectively. Not surprisingly, the use of a constant value to initialise the saliency gives the worst results. This shows that the use of a priori information to define initial saliency improves the vessel extraction quality.

Hessian-based analysis initialisation

Two different initialisation schemes based on the analysis of the Hessian matrix were evaluated: the use of the eigenvalues κ_i of H, and a modified scheme to reflect a token's vesselness. For this second approach, we assigned the response of vesselness filter [10], i.e. $\lambda_1 = \nu(\mathbf{p})$, while the other two λ values were set to zero. The DSC obtained using the response of vesselness filter (0.88 ± 0.04) was higher than the one obtained by directly using the κ_i of the Hessian matrix (0.76 ± 0.05).

Structure tensor-based initialisation

Tokens initialised through the tensorised image gradient were assigned a saliency measurement equal to the squared gradient magnitude $\|\nabla I\|^2$. The value used for the Gaussian kernel size involved in the computation of the structure tensor was obtained from the multi-scale analysis. A DSC of 0.65 ± 0.04 was reported.

Summary

Table 3 summarises the mean Dice score coefficients obtained when comparing the proposed method, using the three different initialisation strategies (no preferred orientation, Hessian-based and structure tensor-based), to the consensus

Table 3 Dice score coefficients obtained using the three different initialisation strategies: no preferred orientation, Hessian-based analysis and structure tensor-based initialisation

	No orientation	Hessian analysis	Structure tensor
DSC	0.89 ± 0.04	0.88 ± 0.04	0.65 ± 0.04

M. For no preferred orientation and Hessian-based initialisation schemes, only the best performing configuration is reported (i.e. using the vesselness filter result). The obtained DSCs using no orientation (0.89 ± 0.04) and Hessian-based initialisation (0.88 ± 0.04) indicate similar performance. Performance using structure tensor initialisation, on the other hand, appears to be worse (0.65 ± 0.04). A visual inspection of the probability maps φ obtained through each strategy (Fig. 2) showed that structure tensor-based initialisation tends to extract big vessels, but fails in the extraction of the small ones, explaining its lower DSC.

Multi-modal versus single-modality extraction

To evaluate the relevance of the multi-modal approach, we compared the proposed method with the results obtained when using a single modality. We evaluate the results obtained when combining the two modalities through:

$$\varphi_{\min}(\mathbf{p}) = \min(s_1, s_2)$$
$$\varphi_{\max}(\mathbf{p}) = \max(s_1, s_2). \qquad (9)$$

It should be noted that the min and max operations are used instead of the logical AND and OR operators, since the images are probabilistic, rather than binary maps.

Results show that when using a single image modality, CTA tends to perform better than 3DPC (Fig. 4). By combining the modalities through the min and max operations, the performance w.r.t 3DPC is improved, but it degrades performance of CTA. To understand this behaviour, we visually inspected the maps obtained through both methods. When using min, weak vesselness responses and discontinuities are propagated into the final map. On the other hand, max favours continuity, but at the cost of increasing the number of false-positives mainly close to boundaries. As a result of this, the DSC can drop dramatically in some cases (see outliers for max in Fig. 4). The proposed method offers a productive balance between the two: False-positives diminish where there is no directionality agreement and weak vesselness responses are not suppressed.

Multi-modal extraction versus single-modality observers

The mean Dice coefficients obtained when comparing our method and the observer's annotations to the consensus M are summarised in Table 4. The DSC of the proposed multi-modal approach is superior to the one obtained by the best performing rater using a single modality. Although CTA images have richer vessel content that is reflected in better rater segmentations, 3DPC contains complementary information that is exploited by the proposed algorithm. A visual comparison of obtained vesselness maps with a consensus map is given in Fig. 3 to further illustrate the performance of our method w.r.t. the current semi-automated approach (Fig. 4).

Computational time

To demonstrate the improvement in computational time due to the modification of the multi-scale analysis, we compared the execution times of the current method and the original one [20] as a function of the number of scales. Both methods were executed on a PC with a quad-core processor (2.13
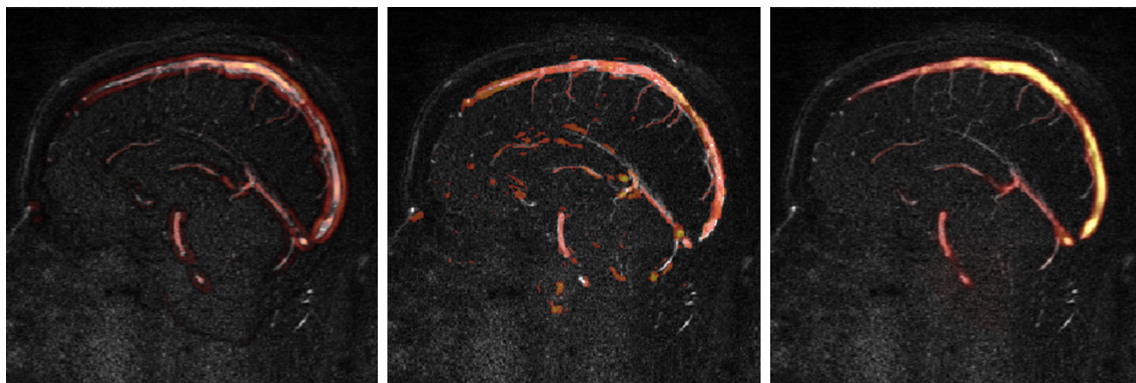

Fig. 2 Probability maps φ obtained using, from *left* to *right*, no preferred orientation, Hessian-based analysis and structure tensor-based initialisation. The response of a vesselness filter [10] was used as initial saliency measurement for the first two cases. The approaches using the response of the vesselness filter are more sensible to fine structures. The structure tensor approach fails to detect small vessels, but has a strong response in large vessels

Table 4 Mean ± standard deviation of the Dice similarity coefficient (DSC) when comparing our method and the observers annotations to the consensus agreement

	Our method	Observer 1		Observer 2		Observer 3	
		3DPC	CTA	3DPC	CTA	3DPC	CTA
DSC	0.89 ± 0.04	0.35 ± 0.10	0.76 ± 0.04	0.37 ± 0.10	0.80 ± 0.03	0.37 ± 0.09	0.79 ± 0.03
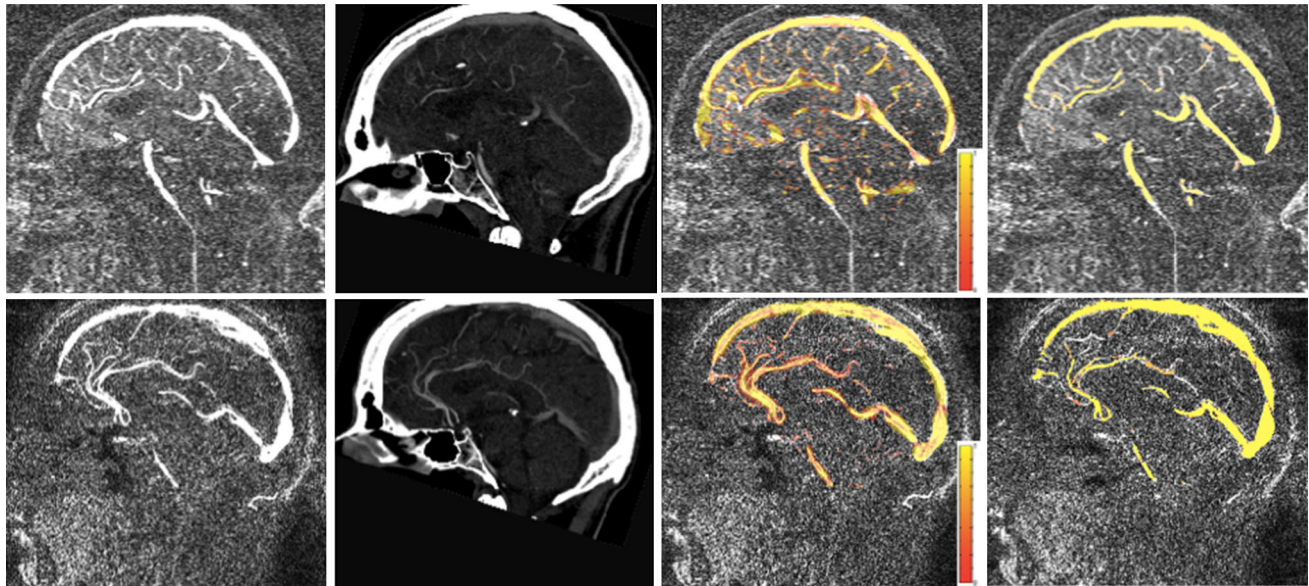
Fig. 3 3DPC (*first column*) and CTA images (*second column*), superposed vesselness map generated by the proposed method over 3DPC (*third column*) and consensus for two subjects (*fourth column*)
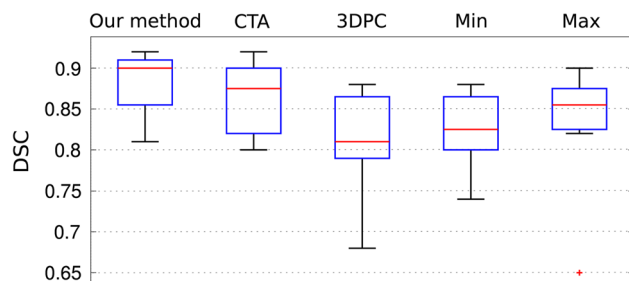
Fig. 4 *Boxplots* displaying the DSC for the proposed method, the single-modality results (without data fusion) using CTA and 3DPC, and data fusion through min and max operators. The *red cross* represents an outlier

GHz). Figure 5 shows the reported execution times along with the Dice score coefficients obtained by each method. Results show that the computational time of the new approach decreases at no cost to the method's accuracy, which remains unchanged.

Integration into the SEEG planning system: visual assessment

Figure 6 presents different examples of the visualisation within the SEEG planning system. Furthermore, we compare the results obtained with the proposed method (Fig. 6 bottom, left) with the extracted vessels when 3DPC and CTA are processed separately (Fig. 6, bottom, right). The images show how the developed method combines the information coming from both modalities while remaining robust to noise.

Discussion

Brain vessels are among the most critical landmarks that need to be assessed in order to mitigate surgical risks [4]. Traditionally, the procedure of extracting vessels (and other structures) is performed in a semi-automated manner by an expert inspecting an image. This is time-consuming, difficult and prone to errors due to the complexity of the vessel network. The development of fully automated and robust methods that can reduce the workload is highly desirable. Here, we have presented a fully automated method that integrates scale, neighbouring structure and feature stability within a single framework to improve vessel extraction within an SEEG planning system.

The proposed method is built upon the tensor voting framework [11]. We have extended it by introducing the evaluation of multiple scales and by using complementary
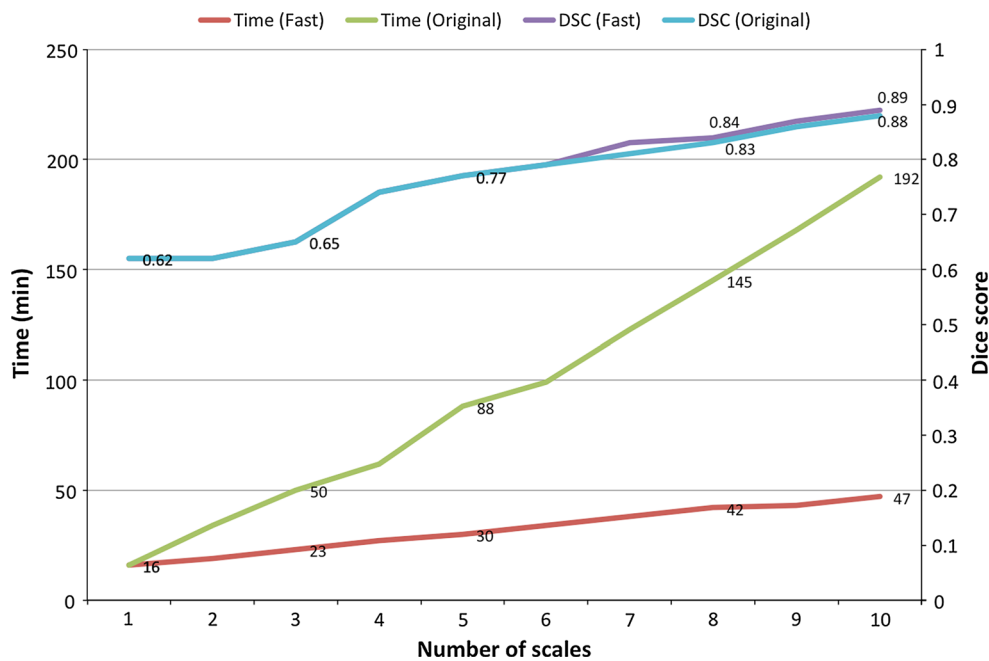
Fig. 5 Average execution times of the proposed approach (fast) and our original formulation Zuluaga et al. [20] as a function of the number of scales. Dice score coefficients (DSCs) for both methods are also displayed to show that speedup of the method is not at the cost of accuracy

sources of information to reduce noisy structures and to improve the connectivity of voxels. Although we have evaluated the proposed framework with two image modalities in this work, its formulation is generic enough that it can be applied to any number of modalities.

The tensor voting framework requires encoding of grey-scale information into a tensorised form. We have evaluated a set of different alternatives for tensor initialisation, which include the use of no priors, zero-order (intensity), first-order (the structure tensor) and second-order information (Hessian-based). Our results have shown that the information derived from the analysis of the Hessian matrix, when used to initialise saliency, provides the best results. Results also show that the selection of optimal saliency measures is more critical than initial tensor orientations. The Dice score coefficients obtained when using vesselness measures as saliency were the highest, independently of the tensor orientation used (no orientation or Hessian-based). On the other hand, changes in the type of saliency measurement used for initialisation did greatly influence the extraction accuracy.

By comparing our method with its single-modality equivalent (e.g. tensor voting without data fusion), we have shown that the use of multiple image types increases performance. This is simply explained by the complementary information offered by the different images. We have also found that the way the information from all sources is combined influences the algorithmic performance. Through the use of a priori knowledge on tubular structures, we favour the fusion of well-aligned objects. Many of the previously seen false-

positives are absent form the final probability map as a result, particularly along the brain boundaries.

It is common to find weak vesselness responses in different regions of the vessel tree in both 3DPC and CTA, and discontinuities are likely to appear there if a single modality is used for visualisation. By combining several modalities, the response of weak vessels increases, improving the continuity of the vessel tree and its visualisation within the planning system. The improvement in the continuity of the vessel tree, combined with a decrease in false-positives detected, should lead to a better path planning.

The comparison of the proposed method with human raters has shown that the use of combined image modalities represents an advantage w.r.t current practice. The presented results are more accurate than a human observer using a single modality. However, the method has some limitations that need to be solved before it can be deployed in clinical practice. While discontinuities are reduced w.r.t single-modality segmentation, they can still appear in small branches. Following the principle of using complementary sources of information (i.e. different image modalities of the same object/anatomical structure), a natural extension of the presented approach to solve this limitation is to combine the outputs obtained when processing different sources of information (e.g. zero-, first- and second-order information or S-,C- and J-maps) into a single probability map. As an example, the structure tensor provides the best results on larger vessels but not performs poorly on small vessels (Fig. 2). Under this scenario, the best features of each source of infor-

Fig. 6 Integration into the computer-assisted planning system. On *top*, examples of displayed extracted vessels using different *colour schemes*. On *bottom left*, display of a segmented vascular tree contrasted with the combined single-modality segmentation, *right*, from 3DPC (*blue*) and CTA (*gold*). The results obtained with the proposed method contain less noise
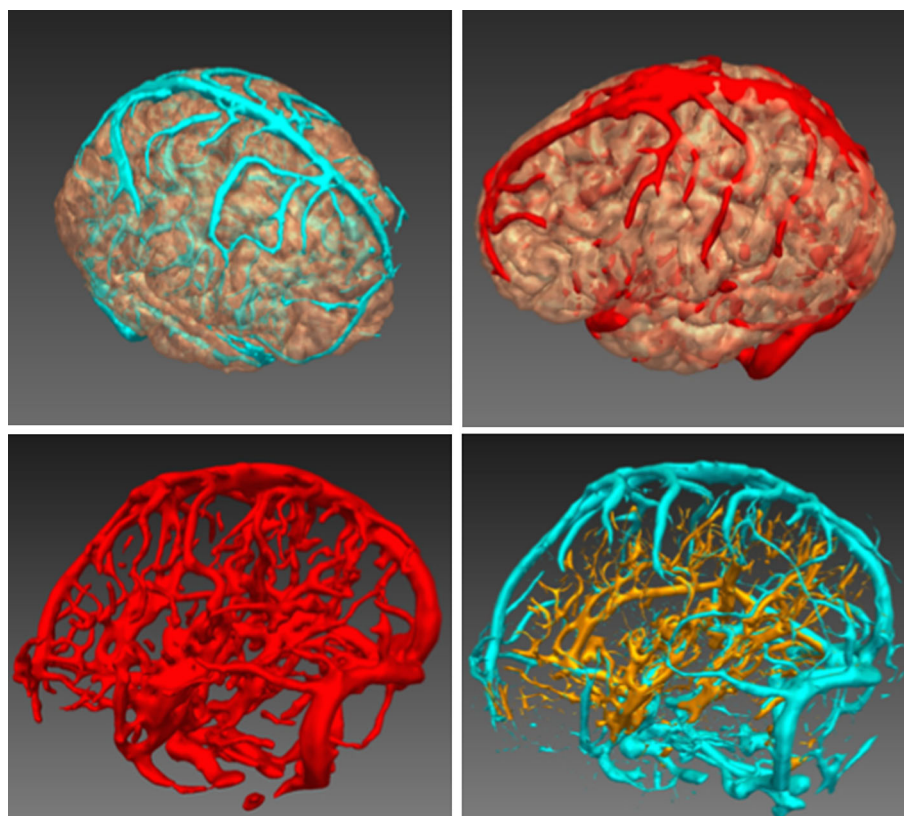
mation would be exploited instead of just keeping the one that overall performs best.

Regarding computational time, which is a key feature if an algorithm is to be translated into the clinic, we have reformulated the multi-scale analysis to reduce the time required to extract the vessel tree. The results show that the new formulation is nearly four times faster than the original one [20] at no cost for the vessel extraction accuracy.

Conclusions

In this paper, we have presented a vessel extraction method for the identification of critical landmarks within a computer-assisted SEEG planning system. The main feature of this method is that it integrates scale, neighbouring structure and feature stability within a single framework. The introduction of a voting neighbourhood within the well-established multi-scale approach and the use of complimentary sources of information reduces noisy structures and improves the connectivity of the voxels belonging to vessels. The results presented here demonstrate the superiority of our method to the semi-automated single-modality segmentation, indicating the possibility of safer SEEG planning with reduced patient morbidity.

Appendix: Stick, plate and ball tensor voting

For the sake of completeness, we present in this section the expressions for SV, PV and BV, the stick, plate and ball votes. Further details on the mathematical derivation of the three expressions can be found in [11,12].

Stick tensor vote

Given $S_{\mathbf{q}}$, which encodes the orientation of the normal at point $\mathbf{q}$, a vote casted from $\mathbf{q}$ to $\mathbf{p}$ can be defined as:

$$\mathrm{SV}(\mathbf{v}, S_{\mathbf{q}}) = s_{1s} \left[R_{2\theta} \; S_{\mathbf{q}} \; R_{2\theta}^T \right], \tag{10}$$

where θ represents the angle defining the orientation of $\mathbf{v}$, $R_{2\theta}$ is a rotation w.r.t the axis $\mathbf{v} \times (S_{\mathbf{q}}\mathbf{v})$, and s_{1s} is a decaying function used to weight the vote and defined as:

$$s_{1s}(\mathbf{v}, S_{\mathbf{q}}) = \begin{cases} e^{-\frac{l^2 + b\kappa^2}{\rho^2}} & : \theta \leq \pi/4 \\ 0 & : \text{otherwise}, \end{cases} \quad (11)$$

where l denotes the arc length, κ the curvature, ρ is a scale parameter associated with the window size $N_{\mathbf{p}}$, and b is a parameter that can be adjusted to give more importance to the curvature ($b = \|\mathbf{v}\|^2/4$, [12]).

Plate tensor vote

Given $S_{P_{\mathbf{q}}}(\beta)$, a stick inside a plate P, a plate vote is defined as the aggregation of stick votes cast by all the stick tensors $S_{P_{\mathbf{q}}}(\beta)$ that constitute $P_{\mathbf{q}}$. It is defined as:

$$PV(\mathbf{v}, P_{\mathbf{q}}) = \frac{\lambda_{1P_{\mathbf{q}}}}{\pi} \int_0^{2\pi} SV(\mathbf{v}, S_{P_{\mathbf{q}}}(\beta)) d\beta. \quad (12)$$

Ball tensor vote

Let $S_B(\phi, \psi)$ be a unitary stick tensor expressed in spherical coordinates with orientation $(1, \phi, \psi)$. Similar to the plate tensor vote, a ball vote can be obtained through the aggregation of the stick votes within the surface Γ of the unitary sphere:

$$BV(\mathbf{v}, B_{\mathbf{q}}) = \frac{3\lambda_{1B_{\mathbf{q}}}}{4\pi} \int_{\Gamma} SV(\mathbf{v}, S_{B_{\mathbf{q}}}(\phi, \psi)) d\Gamma. \quad (13)$$

References

1. Bériault S, Al Subaie F, Collins D, Sadikot A, Pike G (2012) A multi-modal approach to computer-assisted deep brain stimulation trajectory planning. Int J Comput Assist Radiol Surg 7:687–704
2. Cossu M, Cardinale F, Castana L, Citterio A, Francione S, Tassi L, Benabid ALM, Lo Russo G (2005) Stereoelectroencephalography in the presurgical evaluation of focal epilepsy: a retrospective analysis of 215 procedures. Neurosurgery 57(4):706–718
3. David O, Blauwblomme T, Job AS, C S, Hoffmann D, Minotti L, Kahane P (2011) Imaging the seizure onset zone with stereo-electroencephalography. Brain 134:2898–2911
4. Du X, Ding H, Zhou W, Zhang G, Wang G (2013) Cerebrovascular segmentation and planning of depth electrode insertion for epilepsy surgery. Int J Comput Assist Radiol Surg 8:905–916
5. Essert C, Haegelen C, Lalys F, Abadie A, Jannin P (2012) Automatic computation of electrode trajectories for deep brain stimulation: a hybrid symbolic and numerical approach. Int J Comput Assist Radiol Surg 7(4):517–532
6. Frangi AF, Niessen WJ, Vincken KL, Viergever MA (1998) Multiscale vessel enhancement filtering. In: Wells W, Colchester A, Delp SL (eds) Medical Image Computing and Computer-Assisted Intervention—MICCAI, LNCS, vol 1496. Springer, Berlin, pp 130–137
7. Hameeteman K, Zuluaga MA, Freiman M, Joskowicz L, Cuisenaire O, Valencia LF, Gülsün M, Krissian K, Mille J, Wong W, Orkisz M, Tek H, Hoyos MH, Benmansour F, Chung A, Rozie S, van Gils M, van den Borne L, Sosnam J, Berman P, Cohen N, Douek P, Sánchez I, Aissat M, Schaap M, Metz C, Krestin G, van der Lugt A, Niessen W, van Walsum T (2011) Evaluation framework for carotid bifurcation lumen segmentation and stenosis grading. Med Image Anal 15:477–488
8. Hu Z, Niemeijer M, Abràmoff MD, Garvin MK (2012) Multimodal retinal vessel segmentation from spectral-domain optical coherence tomography and fundus photography. IEEE Trans Med Imaging 31:1900–1911
9. Lesage D, Angelini ED, Bloch I, Funka-Lea G (2009) A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. Med Image Anal 13(6):819–845
10. Manniesing R, Viergever MA, Niessen WJ (2006) Vessel enhancing diffusion: a scale space representation of vessel structures. Med Image Anal 10(6):815–825
11. Medioni G, Lee MS, Tang CK (2000) A computational framework for segmentation and grouping. elsevier Science, Amsterdam
12. Moreno R, Garcia MA, Puig D, Pizarro L, Burgeth B, Weickert J (2011) On improving the efficiency of tensor voting. IEEE Trans Patt Anal Mach Intell 33(11):2215–2228
13. Olivier A, Boling W, Tanriverdi T (2012) Techniques in epilepsy surgery: the MNI approach. Cambridge University Press, Cambridge
14. Ourselin S, Roche S, Subsol G, Pennec X, Ayache N (2001) Reconstructing a 3D structure from serial histological sections. Image Vis Comput 19(1–2):25–31
15. Passat N, Ronse C, Baruthio J, Armspach JP, Foucher J (2007) Watershed and multimodal data for brain vessel segmentation: application to the superior sagittal sinus. Image Vis Comput 25:512–521
16. Rosenow F, Lüders H (2001) Presurgical evaluation of epilepsy. Brain 124:1683–1700
17. Sato Y, Nakajima S, Shiraga N, Atsumi H, Yoshida S, Koller T, Gerig G, Kikinis R (1998) Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. Med Image Anal 2(2):143–168
18. Shamir RR, Joskowicz L, Tamir I, Dabool E, Pertman L, Ben-Ami A, Shoshan Y (2012) Reduced risk trajectory planning in image-guided keyhole neurosurgery. Med Phys 39:2885–2895
19. Zombori G, Rodionov R, Nowell M, Zuluaga MA, Clarkson M, Micallef C, Diehl B, Wehner T, Miserochi A, McEvoy A, Duncan J, Ourselin S (2014) A computer assisted planning system for the placement of seeg electrodes in the treatment of epilepsy. Inform Process Comput-Assist Interv Lect Notes Comput Sci 8498:118–127
20. Zuluaga MA, Rodionov R, Nowell M, Achhala S, Zombori G, Cardoso MJ, Miserocchi A, McEvoy AW, Duncan JS, Ourselin S (2014) SEEG trajectory planning: Combining stability, structure and scale in vessel extraction. In: Golland P, Hata N, Barillot C, Hornegger J, Howe R (eds) Medical Image Computing and Computer-Assisted Intervention - MICCAI 2014. Lecture Notes in Computer Science, vol 8674. Springer, Berlin, pp 651–658