Medical Image Analysis 19 (2015) 220–249

Contents lists available at ScienceDirect

Medical Image Analysis

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

Joint 3-D vessel segmentation and centerline extraction using oblique Hough forests with steerable filters

Matthias Schneider [a,b,*], Sven Hirsch [a,c], Bruno Weber [b], Gábor Székely [a], Bjoern H. Menze [a,d]

[a] Computer Vision Laboratory, ETH Zurich, Sternwartstrasse 7, 8092 Zurich, Switzerland
[b] Institute of Pharmacology and Toxicology, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland
[c] Institute of Applied Simulation, ZHAW Wädenswil, 8820 Wädenswil, Switzerland
[d] Institute for Advanced Study and Department of Computer Science, Technische Universität München, Lichtenbergstrasse 2a, 85748 Garching, Germany

ARTICLE INFO

Article history:
Received 13 January 2014
Received in revised form 12 September 2014
Accepted 25 September 2014
Available online 15 October 2014

Keywords:
Vessel segmentation
Centerline extraction
Steerable filters
Oblique random forest
Multivariate Hough voting

ABSTRACT

Contributions: We propose a novel framework for joint 3-D vessel segmentation and centerline extraction. The approach is based on multivariate Hough voting and oblique random forests (RFs) that we learn from noisy annotations. It relies on steerable filters for the efficient computation of local image features at different scales and orientations.
Experiments: We validate both the segmentation performance and the centerline accuracy of our approach both on synthetic vascular data and four 3-D imaging datasets of the rat visual cortex at 700 nm resolution. First, we evaluate the most important structural components of our approach: (1) Orthogonal subspace filtering in comparison to steerable filters that show, qualitatively, similarities to the eigenspace filters learned from local image patches. (2) Standard RF against oblique RF. Second, we compare the overall approach to different state-of-the-art methods for (1) vessel segmentation based on optimally oriented flux (OOF) and the eigenstructure of the Hessian, and (2) centerline extraction based on homotopic skeletonization and geodesic path tracing.
Results: Our experiments reveal the benefit of steerable over eigenspace filters as well as the advantage of oblique split directions over univariate orthogonal splits. We further show that the learning-based approach outperforms different state-of-the-art methods and proves highly accurate and robust with regard to both vessel segmentation and centerline extraction in spite of the high level of label noise in the training data.

1. Introduction

Segmentation and analysis of tubular structures such as blood vessels, in particular, play a crucial role for numerous medically oriented applications and have attracted a lot of attention in the field of medical image processing. The multi-scale nature of vessels, image noise and contrast inhomogeneities make it a challenging task. In this context, a large variety of methods have been developed exploiting photometric and structural properties of tubular structures.

1.1. Related work

1.1.1. Vessel segmentation

Vessel segmentation is an established field in biomedical image processing, see for example Kirbas and Quek (2004) and Lesage et al. (2009) providing extensive reviews. Many of them are notably tailored to specific applications and imaging modalities. Rather simple methods for vessel detection, e.g., absolute or locally adaptive thresholding (Otsu, 1979; Canny, 1983), are regularly used in practice due to their conceptual simplicity and computational efficiency but they are a serious source of error and require careful parameter selection. More sophisticated segmentation techniques can roughly be divided into two groups. One group pursues a top-down strategy by iteratively propagating segmentation labels starting at set of seeds towards distal branches by means of, e.g., region growing (Martínez-Pérez et al., 1999; Lo et al., 2010), active contours (Lorigo et al., 2001), particle filtering (Lesage et al., 2008; Florin et al., 2006), or path tracing (Zhou et al., 2007;

* Corresponding author at: ETF D114.1, Sternwartstrasse 7, 8092 Zurich, Switzerland. Tel.: +41 44 63 20379; fax: +41 44 63 21199.
E-mail address: schneider@vision.ee.ethz.ch (M. Schneider).
URL: http://www.vision.ee.ethz.ch/~schneima (M. Schneider).

http://dx.doi.org/10.1016/j.media.2014.09.007
1361-8415/

Schneider and Sundar, 2010). The design and choice of an appropriate energy or posterior density term to guide the evolution of the segmentation is crucial and usually involves strong assumptions on the underlying structures to be detected. Similarly, elaborate break criteria are required to prevent the segmentation from leaking into the background, particularly for data with a low signal to noise ratio. Another group of methods follows the bottom-up paradigm aiming at optimizing a global neighborhood graph in order to incorporate spatial context (Türetken et al., 2012; Rempfler et al., 2014). The graph is commonly defined on the voxel locations based on a likelihood for a voxel belonging to a tubular structure as well as certain constraints for better robustness, e.g., handling of bifurcations and low-contrast regions (Breitenreicher et al., 2013). Standard optimization strategies such as belief propagation or graph cuts are commonly applied to find the global optimum of the graph which intrinsically defines the termination criteria. However, dealing with large image data, global optimization easily becomes computationally infeasible.

1.1.2. Vessel enhancement

An essential element of all algorithms are measures for tubularity or "vesselness". They are commonly calculated based on optimal filtering and Hessian-based approaches relying on idealized appearance and noise models to enhance tubular structures. The former includes optimal edge detection (Canny, 1983) and steerable filters providing an elegant theory for computationally efficient ridge detection at arbitrary orientations (Jacob and Unser, 2004; González et al., 2009b). The latter is based on the eigenanalysis of the Hessian capturing the second-order structure of local intensity variations (Sato et al., 1997; Frangi et al., 1998). The Hessian is commonly computed by convolving the image patch with the partial second-order derivatives of a Gaussian kernel as the method of choice for noise reduction and to tune the filter response to a specific vessel scale. This basic principle has already been used by Canny for edge and line detection (Canny, 1983; Schneider, 1990). The differential operators involved in the computation of the Hessian are well-posed concepts of linear scale-space theory (Lindeberg, 1996). Modeling vessels as elongated elliptical structures, the eigendecomposition of the Hessian has a geometric interpretation, which can be used to define a vesselness measure as a function of the eigenvalues (Sato et al., 1997; Frangi et al., 1998). Due to the multi-scale nature of vascular structures, Hessian-based filters are commonly applied at different scales. Besides, the eigenvector corresponding to the largest eigenvalue of the Hessian computed at the most discriminative scale is a good estimate for the local vessel direction. In practice, vesselness filters tend to be prone to noise and have difficulty in detecting vessel parts such as bifurcations not complying with the intrinsic idealized appearance model. More recently, Xiao et al. (2013) proposed to replace the Gaussian kernel of standard Hessian approaches with a bi-Gaussian function that allows for independent selection of different scales in the foreground and background. The authors show that a proper selection of the scale parameters reduces interference from adjacent objects while preserving intra-region smoothing. As compared to Hessian-based approaches using inappropriately broad Gaussian kernels, it is hence better suited to resolve neighboring structures, in particular. Vesselness filters have also been successfully applied for global vessel segmentation in X-ray angiography using ridge tracking (Schneider and Sundar, 2010) and graph cut theory (Hernández-Vela et al., 2011). In order to partly overcome the shortcomings of Hessian-based filters, optimally oriented flux (OOF) as introduced by Law and Chung (2008) and its anisotropic variations (Benmansour and Cohen, 2011) have recently gained attention for the segmentation of different anatomical structures including vessels (Benmansour et al., 2013) and intervertebral discs (Law et al., 2013). Briefly, OOF aims at computing an optimal projection direction minimizing the inward oriented flux at the boundary of localized circles (2-D) or spheres (3-D) of different radii (scales). Similar to the Hessian-based approaches, OOF can be used to estimate the local vessel direction as a generalized eigenvalue problem. At the same time, the OOF descriptor is more robust against image noise and local intensity inhomogeneities in the presence of nearby structures, which adversely affects the differential nature of the Hessian. The OOF value, i.e., the projected outward flux, at a certain position and scale can be interpreted as the likelihood of the voxel being centered in a tubular structure of the selected scale. By design, OOF hence provides strong responses at centerlines of curvilinear structures. Similar to the Hessian-based vesselness, the OOF eigenvalues can be combined to obtain a response across the entire structure (Law and Chung, 2008; Benmansour and Cohen, 2011). Finally, Law and Chung (2010) have demonstrated that different measures of image gradient symmetry can be derived from OOF to guide an active contour model for 3-D vessel segmentation with promising results on clinical intracranial and cardiac image data.

1.1.3. Centerline extraction

For many applications, vessel detection, i.e., binary segmentation of the vessel lumen, is insufficient and a more comprehensive vascular description is desirable to characterize the topology and morphology of vascular networks. To this end, the tubular shape of a vascular segment can be modeled by its centerline, i.e., the 1-D curve centered inside the vessel lumen, along with an estimate of the vessel diameter along the centerline assuming a circular cross-section. Other centerline models rely on more general cross-sectional contours such as ellipses (Krissian et al., 2006). Various approaches for centerline extraction have been proposed in the literature including skeletonization by homotopic thinning (Palágyi and Kuba, 1998; Pudney, 1998) and minimal path techniques (Lesage et al., 2009, Section 4.4). The latter computes the centerline between two-points as the path minimizing a certain energetic potential favoring centerline positions. Minimal path techniques enjoy great popularity due to their robustness and global optimality properties (Cohen and Kimmel, 1997). Different variations have been proposed that mostly differ in the definition of the energy term and the numerical optimization schemes such as Dijkstra-like (Gülsün and Tek, 2008; Breitenreicher et al., 2013) and fast marching schemes (Sethian, 1999; Benmansour and Cohen, 2011). Deschamps (2001) defines a distance potential as a non-linear function of the distance to the object boundary. It is used to readjust minimal paths towards the vessel center. Slight inaccuracies in the extracted vessel boundaries may easily impair the distance-based metric, though. Benmansour and Cohen (2011) propose an unisotropic metric based on OOF (Law and Chung, 2008) and present promising results. However, accurate centerline extraction requires a dense sampling of the scale space which is computationally expensive when dealing with tubular structures in a wide range of scales. Recently, voting mechanisms as used for object detection in the computer vision community (Gall et al., 2011) have been applied in the context of centerline extraction to increase robustness against noise and low-contrast regions, in particular (Zhou et al., 2007; Risser et al., 2008; Rouchdy and Cohen, 2012).

1.2. Overview

In this paper, we aim at efficient processing of 3-D high-resolution angiographic image data ($> 10^{10}$ voxels) mapping the cerebrovascular system down to the capillary level, which is of great interest for the analysis of the cerebral vasculature