Computerized Medical Imaging and Graphics 52 (2016) 28–43

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

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

A multi-scale tensor voting approach for small retinal vessel segmentation in high resolution fundus images

Argyrios Christodoulidis[a,*], Thomas Hurtut[a], Houssem Ben Tahar[b], Farida Cheriet[a]

[a] Polytechnique Montréal, Montréal, QC H3C 3A7, Canada
[b] Diagnos Inc., Brossard, QC J4Z 1A7, Canada

ARTICLE INFO

Article history:
Received 24 July 2015
Received in revised form 16 April 2016
Accepted 1 June 2016

Keywords:
Diabetic retinopathy
Fundus imaging
Retinal blood vessel segmentation
Multi-scale line detection
Perceptual organization

ABSTRACT

Segmenting the retinal vessels from fundus images is a prerequisite for many CAD systems for the automatic detection of diabetic retinopathy lesions. So far, research efforts have concentrated mainly on the accurate localization of the large to medium diameter vessels. However, failure to detect the smallest vessels at the segmentation step can lead to false positive lesion detection counts in a subsequent lesion analysis stage. In this study, a new hybrid method for the segmentation of the smallest vessels is proposed. Line detection and perceptual organization techniques are combined in a multi-scale scheme. Small vessels are reconstructed from the perceptual-based approach via tracking and pixel painting.

The segmentation was validated in a high resolution fundus image database including healthy and diabetic subjects using pixel-based as well as perceptual-based measures. The proposed method achieves 85.06% sensitivity rate, while the original multi-scale line detection method achieves 81.06% sensitivity rate for the corresponding images ($p < 0.05$). The improvement in the sensitivity rate for the database is 6.47% when only the smallest vessels are considered ($p < 0.05$). For the perceptual-based measure, the proposed method improves the detection of the vasculature by 7.8% against the original multi-scale line detection method ($p < 0.05$).

1. Introduction

Diabetic retinopathy (DR) is a public health issue with major implications for the eyesight of the affected population. In diabetes, blood glucose levels are higher than in the normal case. Prolonged hyperglycemia has a deteriorating effect upon retinal circulation that triggers many pathophysiological changes, of which the most important are related to the smallest vessels (Archer, 1999). If it is left untreated, this disease can progressively lead to blindness. The prevalence of blindness caused by diabetic retinopathy is 4.8% out of the total of 37 million cases throughout the world (World Health Organisation, 2005). The introduction of screening programs has helped to reduce the incidence of diabetic retinopathy (Scanlon, 2008). However, a major drawback of current screening techniques is their reliance upon visual inspection of fundus images by physicians. Computerized systems assisting retinal specialists' decisions, known as computer-aided diagnosis (CAD) systems, would allow screening programs to cover larger populations and reduce the disease's impact though earlier detection.

Segmentation algorithms can highlight the most important retinal structures and guide the physicians to quickly identify signs of DR. A major category of CAD systems for DR rely on pre-segmenting structures other than dark or bright lesions, i.e. the vessels and the optic disc, before proceeding to the lesion analysis stage (Abràmoff et al., 2010; Antal and Hajdu, 2013; Sopharak et al., 2013). However, the vasculature tends to be coarsely extracted and thus among the candidate lesions, many small and possibly fragmented vessels exist. Consequently, the system can mistake these fragments for dark lesions because they share common features with the latter (Sopharak et al., 2013). To this end, the focus of this work is the segmentation of small retinal vessels that can potentially improve an automatic diagnosis system for retinopathies.

Despite the substantial amount of research on retinal vessel segmentation (Abràmoff et al., 2010; Fraz et al., 2012; Winder et al., 2009), the algorithms that have been proposed focus on the larger vessels while omitting the smaller ones. Since the latter are thinner, with lower contrast than average vessels and possible discontinuities, they need special handling in a separate step. In addition, the performance assessment for these methods is based solely on global binary classification metrics (Fawcett, 2006) that

* Corresponding author.
E-mail addresses: argyrios.christodoulidis@polymtl.ca (A. Christodoulidis), thomas.hurtut@polymtl.ca (T. Hurtut), hbentahar@diagnos.ca (H.B. Tahar), farida.cheriet@polymtl.ca (F. Cheriet).

http://dx.doi.org/10.1016/j.compmedimag.2016.06.001
0895-6111/

do not take into account the fact that the smallest vessels represent approximately 10% of the total surface area of the vessel network (Niemeijer et al., 2004).

To our knowledge, only a limited number of existing studies concentrate on small vessel segmentation in high resolution images, with a reliable evaluation against manual segmentation. The available methods can be grouped in two categories: supervised machine learning and intensity-based approaches. Machine learning approaches (Kovács and Hajdu, 2011; Sofka and Stewart, 2006) require an a priori model trained on a labeled set. This set is usually small because the labeling task is laborious, and not representative of the variability observed in the real world. Even though small and large vessels share common morphological attributes such as elongateness or darker appearance than their background, different kinds of features must be selected in order to perform classification, and the detected features training should be confined to the smallest vessel category. This was partially met in Kovács and Hajdu (2011) where the training phase was divided into large and small vessels; however, during the classification stage both categories were merged and the segmentation performance was assessed globally, without examining the small vessels separately. This method has been used for vasculature subtraction prior to microaneurysm detection in Antal and Hajdu (2013). Sofka and Stewart (2006) proposed a new vesselness measure that is robust to lesions away from the vasculature. The method was validated on small and poorly contrasted vessels, but its output is only the vessel centerlines, which is inadequate if the result must be directly used for lesion detection in a CAD system. In that case, full vasculature segmentation must be performed before any subsequent lesion detection stage.

In the intensity-based methods, a template is applied across the image in the spatial domain in order to segment either full vessels (Zhang et al., 2009; Allen et al., 2011; Oost et al., 2010) or centerlines (You et al., 2011; Mendonca and Campilho, 2006; Xu and Luo, 2010). Although the method in Zhang et al. (2009) proposes an approach to suppress lesion detection away from the vasculature, it is assumed that the small vessels are piecewise straight (Zhang et al., 2009; You et al., 2011; Allen et al., 2011). By contrast, curve fitting was used in Oost et al. (2010) to approximate small vessels' curvatures and to reconnect pre-segmented fragments into larger segments. Another shortcoming of the available intensity-based methods is that the segmentation results are compared quantitatively in all the available categories of vessels without making a distinction between the small and large vasculatures. Moreover, a limited number of methods have been used in the context of CAD systems for lesion detection (Zhang et al., 2010; García et al., 2008). A hybrid method was proposed by Xu and Luo (2010). Small vessel centerlines are first discriminated from noise by classification, then they are merged into larger segments based on local tracking.

Several multi-scale vesselness pipelines have also been proposed to segment the vessels in high resolution images and tested on Erlangen database (Odstrcilik et al., 2013). These works are based on the application of Frangi's model (Frangi et al., 1998) or Gaussian matched filtering (Odstrcilik et al., 2013). These studies modified Frangi's method either by incorporating a pre/post-processing step (Yu et al., 2012; Annunziata et al., 2015; Budai et al., 2013) or by adapting the basic model to handle crossings and bifurcations (Hannink et al., 2014). A Hessian matrix is approximated by partial second-order Gaussian derivatives, but this has the effect of filtering out the small vessels, magnifying the background noise and reconnecting the noise to the detected vessels. Equally, matched filtering is not sensitive enough for the smallest vessels because the template used is too general to represent vessels that demonstrate high curvatures over short distances and limited numbers of cross-sectional vessel entries.

Line detectors are a recent family of methods for identifying linear structures. These were first used to identify linear structures in mammographic data (Zwiggelaar et al., 2004), and recently adapted to the task of detecting retinal vessels (Ricci and Perfetti, 2007). Recently, Nguyen et al. (2013) introduced the multi-scale line detector (MSLD) to accommodate for the various vessel sizes that appear in the retina. This method utilizes a straight sampling segment of variable length that, when rotated around a central pixel, gives the line response. The values of a sliding averaging window are subtracted from the line response to compensate for uneven background illumination. Finally, the line detector responses are thresholded in order to segment the vessels.

The MSLD method efficiently segments large to medium-sized vessels, however several factors hamper the segmentation of the smallest vessels. First, the size of the sliding window is set globally. In the case of a small vessel the window is too large and there are not enough vessel pixels to compensate for the presence of many background values. Second, the small vessels usually appear near the macula and the periphery where the image intensity and contrast are low; as a result, the line response is low and these vessels are fragmented.

In this context, it is useful to consider the tensor voting framework (TVF) developed by Medioni and Kang (2004) and Medioni et al. (2000). This is a bottom-up approach for organizing neighborhood information based on perceptual principles from Gestalt theory used to segment thin fragmented structures. It formalizes the observation that high-level perceptual structures can be formed by grouping individual lower level structures. Tensor voting consists of two components, given as two criteria: (1) the proximity criterion, namely pixels belonging to the same structure are close to each other; and (2) the continuity criterion, namely grouped structures have constant curvature. The process starts by initializing the directions of the structures via a method to estimate orientations, such as ball voting or gradients. Stick voting is then applied to refine the directions, propagate and reconnect the information between the structures, and compute their saliency. Saliency measures the certainty about the existence of a line in a particular location, and it demonstrates a local maximum along the principal direction of the structure. A scale parameter σ_{TVF} controls the range of influence of the voting.

This method has been successfully applied to a range of biomedical imaging applications (Risser et al., 2008; Loss et al., 2011; Yigitsoy and Navab, 2013; Zuluaga et al., 2015; Maggiori et al., 2014). These include filling microvascular discontinuities (Risser et al., 2008) based on interpolating the width between reconnected segments, and detecting important cellular structures (Loss et al., 2011) by iterating a TVF along a range of scales. A variety of algorithmic improvements to the basic framework were proposed in the literature; these were comprehensively reviewed by Maggiori et al. (2014).

In retinal imaging applications, single scale TVF has been applied to either extract full vessels (Park et al., 2007) or segment centerlines (Leng et al., 2011). No particular care was given to the problem of full small vessel segmentation in either of these studies. The work in Park et al. (2007) focused on identifying the optic disc as the converging point of the largest vessels and used a direct thresholding of the saliency map. In this manner, disconnected curvilinear structures belonging to lesions or irrelevant structures can get high saliency values, so thresholding is likely to remove small vessels without suppressing the false positives. The authors in Leng et al. (2011) extracted the centerlines based on steerable tensor voting (Franken et al., 2006). Rather than applying the framework for extracting the saliency of the structures globally, they used the method to link pre-identified edges. Moreover, in a preliminary phase, edge points were discarded based on a global length criterion. But the setting of a non-adaptive global length