Computerized Medical Imaging and Graphics 38 (2014) 49–56

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

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

Approximate Nearest Neighbour Field based Optic Disk Detection

S. Avinash Ramakanth, R. Venkatesh Babu*

Video Analytics Lab, Supercomputer Education and Research Centre, Indian Institute of Science, Bangalore 560 012, India

ARTICLE INFO

Article history:
Received 10 July 2013
Received in revised form 25 October 2013
Accepted 29 October 2013

Keywords:
Approximate Nearest Neighbour Field
FeatureMatch
Optic disk detection
Fundus or retinal image processing

ABSTRACT

Approximate Nearest Neighbour Field maps are commonly used by computer vision and graphics community to deal with problems like image completion, retargetting, denoising, etc. In this paper, we extend the scope of usage of ANNF maps to medical image analysis, more specifically to optic disk detection in retinal images. In the analysis of retinal images, optic disk detection plays an important role since it simplifies the segmentation of optic disk and other retinal structures. The proposed approach uses FeatureMatch, an ANNF algorithm, to find the correspondence between a chosen optic disk reference image and any given query image. This correspondence provides a distribution of patches in the query image that are closest to patches in the reference image. The likelihood map obtained from the distribution of patches in query image is used for optic disk detection. The proposed approach is evaluated on five publicly available DIARETDB0, DIARETDB1, DRIVE, STARE and MESSIDOR databases, with total of 1540 images. We show, experimentally, that our proposed approach achieves an average detection accuracy of 99% and an average computation time of 0.2 s per image.

1. Introduction

In the context of image processing, Approximate Nearest Neighbour Field (ANNF) denotes the correspondence between a pair of images, say ‘source’ and ‘target’, such that every patch in target image maps to a corresponding patch in source image, if it minimises the Euclidean distance between them. These ANNF maps are used widely by the graphics and computer vision community to deal with problems like image completion [1], retargetting [2], denoising [3], optic disk detection [4], etc. One of the major reasons for the extensive use of ANNF maps is the development of efficient algorithms to compute ANNF maps, such as PatchMatch [2] and FeatureMatch [5], using which the correspondence between two HD images (2048 × 1536 pixels) can be computed within a few seconds (<15 s).

In this paper, we propose a novel application of ANNF maps such that, using a single optic disk image as reference dictionary, we can analyse retinal images to detect the optic disk (OD). Retinal image analysis is an important problem for ophthalmologists, since it plays a major role in diagnosis of various diseases like diabetes mellitus, glaucoma, hydrocephalus, benign intracranial hypertension and brain tumour, among other conditions. As studied by Singer et al. [6], nearly 30,000 individuals per million total population require retinal screening for diabetic retinopathy on a regular basis, and nearly 50% reduction in blindness, due to diabetic retinopathy, can be observed by timely screening.

During analysis of retinal images, detection of optic disk is an important pre-processing stage, since it helps to localise and segment other anatomical structures such as fovea (where the distance between OD centre and centre of fovea is roughly constant) [7]. When detecting exudates and lesions, OD can be confused with other lesions, hence localisation of OD is important to remove it from the set of candidate lesions [8]. It also helps in computing important diagnostic indices for hypertensive retinopathy based upon vasculature, such as central retinal artery equivalent (CRAE) and central retinal vein equivalent (CRVE) [9]. Furthermore, it can also be used to classify left and right eyes in fovea-centred retinal images [10].

Conventional methods for the localisation of optic disk are based on two major assumptions: (i) OD is the brightest region of retinal image [9,11–14] or (ii) OD is the point of origination of retinal vasculatures [7,8,15–17]. Though the OD is prominently brighter than surroundings in normal retinal images, the shape, colour and brightness show large variation in the presence of diseases. Due to this, OD detection methods based on retinal vessels are used in place of brightness based methods. Though vessel based OD detection has been shown to perform better than colour based OD detection, there exist few short-comings like parameter tuning and extensive pre-processing for vessel based OD detection. A much more detailed explanation of these existing methods is given in the next section.

* Corresponding author. Tel.: +91 80 22932900; fax: +91 80 23602648.
E-mail addresses: avinashrs@ssl.serc.iisc.in (S.A. Ramakanth), venky@serc.iisc.in (R.V. Babu).
URL: http://val.serc.iisc.ernet.in (R.V. Babu).

0895-6111/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compmedimag.2013.10.007

This paper is organised as follows: we begin with related work in Section 2, followed by the proposed approach in Section 3. Experimental results, benchmarking the performance and computational time of the proposed approach are discussed in Section 4. Finally, Section 5 concludes the paper.

2. Related work

The optic disk, in healthy retinal images, is characterised by a bright circular region [15] from where the optic nerves, retinal and choroid vessels emerge into the eye [7]. Simple methods like searching for brightest region [11] or searching for region with highest image variation [12] work well on healthy retinal images, but fail to localise the OD correctly when retinal images suffer from anomalies, like exudates, which have similar appearance properties as the OD.

To search for the optic disk in the brightest region, Walter et al. [11], used a simple thresholding scheme to obtain a map with all the bright regions in a retinal image. This binary map is used to localise the optic disk as the centre of the biggest and brightest connected component.

The assumption of OD being the brightest region was further used by Li et al. [9,13] in order to find the OD candidate regions in their model-based approach. The intensity image is thresholded such that the brightest 1% pixels are selected as the OD candidate regions. These regions are clustered and small clusters are discarded. An optic disk model is created by applying principal component analysis (PCA) to a training set of 10 intensity normalised square sub-images manually cropped around the OD. For each pixel in the candidate region, the PCA transform is applied over different scales (0.8–1.2) and the OD is localised as the region with the smallest Euclidean distance to its projection onto the disc-space.

Based on the same assumption that OD is a circular bright region, Lu et al. [14] used a line operator to capture the circular brightness of the optic disk. This approach evaluated the image brightness variation along multiple line segments of specific orientations that pass through each retinal image pixel. The orientation of the line segment with the minimum/maximum variation has specific pattern that is used to locate the OD accurately. In a similar approach, Lu [18] used a circular transformation, instead of line transform, to capture both the circular shape of the OD and the image variation across the OD boundary simultaneously.

The shape, colour and size of the OD show large variation especially in the presence of retinopathies, and therefore, detection methods based on these properties were observed to be weak, and impractical. Apart from being a bright circular region, the OD is the entrance point for both the optic nerve, and the main blood vessels. As a result, many techniques have tried to utilise the information provided by the retinal vasculature. In order to determine the origination of the retinal vasculature (convergence point), and thus localise the OD, Hoover et al. used Fuzzy convergence [15] to vote the probable OD locations. The inputs to the fuzzy convergence algorithm are six binary vessel segmentations (each at a different scale) obtained from the green-channel image. Each vessel is modelled by a fuzzy segment, that contributes to a cumulative voting image (a convergence image) where each pixel equals the amount of fuzzy segments on which the pixel lie. Finally, the convergence image is smoothed and thresholded to determine the strongest point(s) of convergence. If the final result is inconclusive, the green-channel is illumination equalised, and Fisher's linear discriminant is applied to regions containing the brightest pixels to detect the OD.

Another model-based (template matching) approach was proposed by Youssif et al. [7]. The algorithm works on direction of retinal vessels and to segment the retinal vessels, 2D Gaussian matched filters are used. The segmented vessels are then thinned, and filtered using local intensity, to represent the OD-centre candidates. A vessel's direction matched filter is used, whose size is specifically tuned to suit various data-sets. The difference between matched filter (re-sized into four different sizes), and the vessels' directions at the surrounding area of each of the OD-centre candidate is measured. The minimum difference provides an estimate of the OD-centre coordinates.

Closely related to vasculature fitting using a directional model, Foracchia et al. [8] identified the position of the OD using a geometrical model on the vessel structure. In this method, the main vessels originating from the OD are geometrically modelled using two parabolas. Consequently, the OD position is located as the common vertex of the two parabolas (i.e., the convergence point for the retinal vasculature). After the complete model of vessels direction is created, the vasculature of the input image is extracted, and the difference between the model directions and the extracted vessels directions is minimised using the weighted residual sum of squares (RSS) and a simulated annealing (SA) optimisation algorithm.

One more model based approach was proposed by Yu et al. [16] which estimates the OD candidate location by matching a template, designed to adapt to different image resolutions. Then, vessel characteristics on the OD are used to determine the final OD location. Initialised by the detected OD centre and estimated OD radius, a fast, hybrid level-set model, which combines region and local gradient information, is applied for the segmentation of the disk boundary. Morphological filtering is used to remove blood vessels and bright regions other than the OD that affect segmentation in the peripapillary region.

The common element in the above template based methods is the extraction of retinal vessels from a retinal image, which is computationally intensive. Furthermore, these methods are based on finely tuned parameters to detect and threshold veins from a retinal image to obtain the vessel map. The parameters in these approaches are tuned for specific data-sets (STARE data-set in [7,8], MESSIDOR data-set in [16]), and to obtain accurate detections on unseen data-sets these parameters will have to be re-tuned for the algorithms to function efficiently. As an improvement over these model based approaches Mahfouz et al. [17] proposed a fast localisation technique, based on projections of image features that encode the x and y gradients of the OD. The resulting 1D projections are then searched to determine the candidate OD locations. This avoids searching the 2D image space and, thus, enhances the speed of the OD localisation process. Image features based on retinal vessels orientation and the OD brightness are used to localise the final OD location.

Apart from these approaches, which model the characteristic features of the optic disk for localisation, Sinha et al. [19] posed the OD detection problem as a classification problem and used a scale-embedded dictionary in conjunction with l_1 minimisation technique to detect the optic disk. To create a dictionary, manually marked fixed-sized sub-images that contain OD at the centre are extracted at multiple scales. For a given test image, all sub-images are sparsely represented as a linear combination of OD dictionary elements. A confidence measure indicating the likelihood of the presence of OD is obtained from these coefficients. Red channel and grey intensity images are processed independently, and their respective confidence measures are fused to form a confidence map. A blob detector is run on the confidence map, whose peak response is considered to be at the location of the OD.

A common feature in all the approaches discussed above is that a single detection scheme is used to localise the OD location. Instead of using a single detection scheme, Qureshi et al. [20] proposed combining multiple OD detection algorithms along with macula detection methods to localise the optic disk in retinal images. Based on the property that an ensemble of algorithms performs better