LOCATING THE OPTIC DISK IN RETINAL IMAGES VIA PLAUSIBLE DETECTION AND CONSTRAINT SATISFACTION

Emanuele Trucco and Pawan J. Kamat

Heriot Watt University, Riccarton, Edinburgh, EH14 4AS, UK
e.trucco@hw.ac.uk

ABSTRACT

This paper present a novel, robust approach to the automatic location of the optic disk in retinal (fundus) images. Instead of generating a single, high-confidence optic disk candidate, we generate sets of *plausible* candidates for optic disk, macula and main vessels, then search the space of all possible triplets (optic disk, macula, vessels) to identify the one satisfying a-priori anatomical constraints. Our first implementation achieved 100% success with 40 wide-field-of-view retinal images acquired by an OPTOS Panorama ophtalmoscope. It also matched the performance of a visible, recently reported algorithm [1] on the STARE test set [2], and succeeded in some cases where [1] failed.

1. INTRODUCTION AND RELATED WORK

This paper present a novel, robust approach to the automatic location of the optic disk (henceforth OD), and more generally of OD, macula, and the main blood vessels (arcades), highlighted in Figure 2. It has been argued that locating such landmarks reliably would facilitate the automatic detection of disease symptoms in retinal images [3, 4], most notably glaucoma, diabethic retinopathy, and age-related macula degeneration. Diabetis, for example, has an incidence of 3 to 5% in the UK, peaking around 20% in some US areas.

A plethora of algorithms have been proposed for the detection of the optic disk and the vasculature in fundus images, and, to a lesser extent, of the macula [3, 5, 6, 1, 7, 8, 4, 9], both *per se* and in the context of a more general analysis; see, e.g., [1] for a recent, concise review focussing on OD. A serious difficulty is that the appearance of the landmarks in question can vary significantly with diseases, imaging conditions (even with the same equipment), race, and other factors. Algorithms relying only on brightness levels are unlikely to prove sufficiently general and reliable.

The majority of the reported approaches strive to detect *individual* landmarks with the highest confidence possible. Confidence is typically estimated by the percentage of true positives (correct detections) in controlled experiments with ground-truth (or "gold standard") available.

We propose an alternative approach taking advantage of the significant a-priori knowledge available on retina anatomy. Instead of a detector expected to identify *the* right OD location with very high confidence, we deploy simple, individual detectors of OD, macula and arcades generating sets of *plausible* candidates for each landmark, among which the true ones are highly probably included. We then search the space of all possible triplets (OD, macula, arcade) formed by the candidates, to identify the one satisfying constraints imposed by retinal anatomy.

This approach offers several advantages. First and most important, the combination of a-priori knowledge and brightness-based analysis combats effectively the uncertainty of brightness-only techniques caused by the variability of fundus images. Searching the right landmark triplet in a potentially large space of plausible solutions bears excellent promises of really robust detection; our first prototype, even with simple detectors and constraints, performs as well as, and on some images better than, Hoover and Goldbaum's recently published algorithm [1] on the challenging STARE test set [2]. Second, at a parity of confidence, it is easier to identify a *group* of several, plausible candidates containing the true landmark (e.g., OD), than to identify only *the* true one; the landmark detectors can therefore be simpler and more efficient. In addition, we extract only the *approximate* path of the main blood vessels, as this is sufficient to impose effective anatomical constraints. The effort required by this is modest if compared to the one needed to detect the actual vessel pixels.

Although constraints inspired to retina anatomy have been used before to detect the OD (e.g., convergence or spatial density of blood vessels), no author has, to our best knowledge, taken advantage of simultaneous constraints on OD, macula and arcades to search a potentially large space of plausible candidates in a full constraint satisfaction framework.

The remainder of this paper describes briefly our current system and its key modules (Section 2) and some experi-

Thanks to OPTOS plc for making their data set available, and to the STARE team for making data and results public.

0-7803-8554-3/04/$20.00 ©2004 IEEE.

mental results with two sets of retinal images (Section 3). We discuss our work and its future in Section 4.

2. THE ALGORITHM

2.1. Defining "detection"

It is necessary to define "detection" as this may be taken to mean different operations. For the OD (resp. macula), we aim to specify a point within the image region that an expert observer would identify as the OD (resp. macula). We are not concerned with shape measurement. For the arcades, we aim to identify the approximate path of these main blood vessels by simply fitting a parabola with vertex near the centre of the relevant OD candidate region. This is because we are *not* interested in the exact vessel pixels, but only in sufficient information to constrain the location of OD and macula.

2.2. Outline of the algorithm

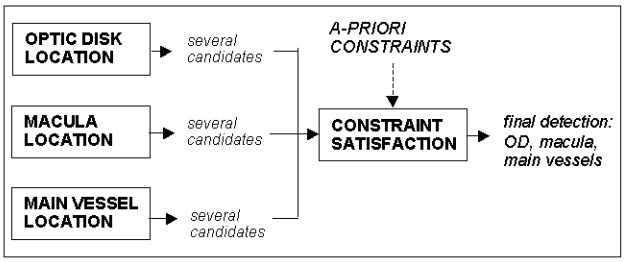

Fig. 1. Basic system architecture. The input image is fed simultaneously to all the location modules.

The architecture of the algorithm is shown in Figure 1. Each block is briefly described below.

2.2.1. Optic disk detection

This simple detector assumes that the OD is one of the brightest image regions [8], although not necessarily the brightest. First, we apply a 7×7 median filter to attenuate image noise. We then apply morphological closing followed by opening to suppress most of the vasculature information. The structural element is a disk larger than the largest vessel cross-section (8-pixel diameter in our case). Finally, we collect the N_o brightest regions in the image (based on mean region intensity) as plausible OD candidates. These may include diseased areas and possible bright, noisy spots.

2.2.2. Macula detection

Similarly to the OD detector, we assume that the macula corresponds to one of the darkest image regions. However, unlike the OD, grey levels in the macula region cannot be expected to be significantly darker than the rest of the image, e.g., than major vessels. We use the same, simple detector adopted for the OD, but keep the N_m darkest candidates. These include, in general, noisy areas (e.g., peripheral low-intensity spots), diseased areas, nevus, and so forth.

2.2.3. Arcade detection

Arcade detection works in three stages.

First, we enhance strong vessels to obtain an image I_e. This is done by subtracting a background image, I_b, from a vessel-enhanced one, I_v, so that $I_e = I_v - I_b$. I_b is computed by morphological close/open, similarly to above; I_v as $I_v = I_t - I_b$, where I_t and I_b are, respectively, the top-hat (difference between the image and its opening) and bottom-hat (difference between the image and its closing) transformations of the original image. The structural element (disk) here is larger than for OD and macula (apprx. 15 pixels diameter).

Second, we compute the edge strength (gradient magnitude) at all pixels of I_e. This results in high values at vessel boundaries, especially so for the typically thick and dark arcades.

Third, we fit a parabola, $C(a, b, c)$, given by $x = ay^2 + by + c$, with axis parallel to the horizontal image axis. We initialize the vertex to be at the centre of the OD candidate region. As much of the correct parabola is expected to overlap the arcades, we maximize a merit function integrating local vessel evidence (gradient magnitude) along the candidate parabola:

$$\max_{(a,b,c)} \sum_{C(a,b,c)} \sum_{W(\mathbf{p}_C)} \|\nabla I_e(\mathbf{p}_C)\|^2,$$

where $W(\mathbf{p}_C)$ is a 7×7 window centered in $\mathbf{p}_C$, the generic point on C. We use all possible non-overlapping windows along the hypothesized curve. Maximization is currently performed by Nelder-Mead simplex search [10].

2.2.4. Constraint satisfaction

We now search the space of the triplets (OD, macula, arcade) generated by combining all plausible candidates for the one consistent with retinal anatomy constraints. This is done with an *interpretation tree* [11] (IT). In essence, an IT considers incrementally growing tuples of candidates, checking each new tuple against pre-defined constraints. If the check fails, the search backtracks by replacing the latest addition. The search terminates successfully when the first globally consistent solution is found. If no globally consistent solution is found, the search fails. There may occasionally be multiple acceptable solutions (e.g., if two OD candidates are very close to each other and both within the

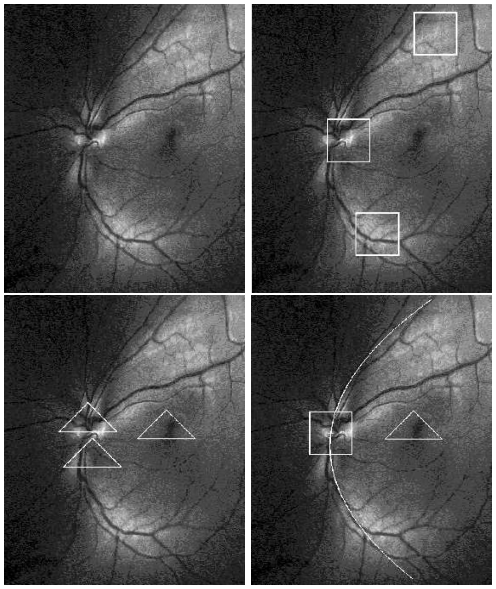

Fig. 2. Top left: input image. Top right: OD candidates. Bottom left: macula candidates. Bottom right: final result.

real OD region), in which case any of them (e.g., the first one) gives a valid detection. Therefore, unlike model-based 2D-3D feature matching [11], there seems to be no need for verification. The triplet is built by exploring OD candidates first, then arcade ones, then macula ones. This ordering reduces the actual space searched, as the OD detector is more reliable than the macula detector. Search is further limited by exploring first the most plausible candidates (for macula, the darker the more plausible; for OD, the brighter the more plausible). Notice that this scheme supports a much higher number of constraints, candidates and objects than the ones used in the tests reported here.

We use currently only four simple constraints:

1. *The macula must be near the approximate axis of symmetry of the arcade*. This is checked by measuring the distance between the hypothesized macula and the axis of the fitted parabola. The constraint fails if $d > \tau_m$, with τ_m a pre-defined threshold depending on the type (size, resolution) of images examined.

2. *The arcade must converge in the OD*. This is implemented by initializing the vertex of the parabola fit to be at the hypothesized OD location.

3. *The distance between OD and macula must fall within a pre-defined interval*. The latter depends again on the type of images examined.

4. *The average constrast along the hypothesized parabola must be high*. This allows us to drop wrong OD candidates placed at acceptable distances from macula candidates, but producing poor arcade fits.

Although image-dependent threshold are best avoided (but similar thresholds feature in other systems, e.g., [1]), we stress that the main contribution of this work is the general paradigm, not the specific components adopted.

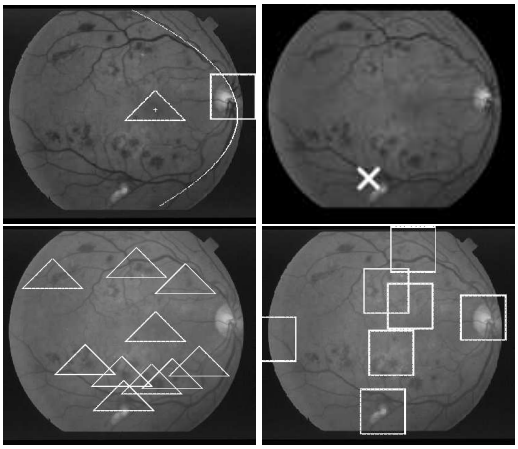

Fig. 3. Top left: correct detection by our algorithms (STARE img0139). Top right: incorrect STARE result. Bottom left: plausible macula candidates. Bottom right: plausible OD candidates.

3. EXPERIMENTAL RESULTS

An initial prototype has been implemented in MATLAB on a 2.4GHz PC under XP. The diameter of the disk used as structural element in the morphology-based OD and macula detection was 8 pixels throughout.

The algorithm produced 100% correct OD detections on 40 wide-FOV green-channel 1984×1984 images of healthy and mildly diseased eyes acquired by an OPTOS Panorama system. We used the 840×700 central part of the images, which is guaranteed to contain all landmarks. τ_m (OD-macula distance) was approximately 250 pixels. The algorithm proved stable (unchanged detection results) in a large range of τ_m (about 100 pixels, i.e., 40% variations). Three candidates per detector ($N_o = N_m = 3$) captured the true OD and macula, giving perfect detection for this set (see below). Figure 2 shows a representative input image, the associated OD and macula candidates, and the final landmark detection.

We also tested our system on the STARE test set of 605×700 images [2], which includes results of Hoover's and Goldbaum recent OD detection algorithm [1]. Given the assumptions of our system, we considered only the 63 images capturing a region including both macula and OD, even if not necessarily visible. This subset included several

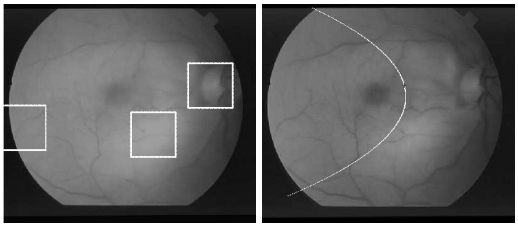

Fig. 4. Incorrect detection by our algorithm but [1] success (STARE img 004). Left: plausible OD candidates (incl. true one). Right: wrong parabola fit for true OD candidate.

retinas showing severe disease symptoms. A maximum of 10 candidates per detector were allowed for OD and macula ($N_o = N_m = 10$). Notice that more candidates were needed to achieve good results than with the previous set; we have not quantified yet the impact of N_o and N_m variations on robustness. Even with the simple detectors used currently, we matched the detection rate of [1] (86% success, or 9 failures) on this difficult subset. Most interestingly, failure is not always common: [1] fails and we succeed on images 7, 19, 27, 139, and vice versa on 4, 45, 219, 240. There are five common failures (3, 13, 20, 26, 41). Figure 3 shows an example where we succeed and [1] fails. Here, our algorithm is not distracted by the bright area on the left, as multiple plausible candidates are included. Figure 4 refers to a case where [1] succeeds and we fail. Notice that the true OD *is* included in the set of plausible candidates, but fails to be selected due to insufficient local vessel support (wrong parabola fit).

4. CONCLUSIONS

We have presented a novel approach to the detection of the OD in retinal images, and more generally of retinal landmarks (OD, macula, and approximate path of the main vessels). Our approach is based on plausible (limited confidence) detection and constraint satisfaction. The strength of the approach is the combination of less stringent demands on individual landmark detectors (which are allowed to return groups of plausible candidates, not just the right one), and a constraint satisfaction mechanism enforcing mutual constraints on all landmarks. Global analysis of retinal landmarks, in addition, identifies basic but important properties of the eye, e.g., left or right. Current results suggest excellent potential for robustness.

Part of our current work is devoted to a full comparative evaluation of our prototype with badly damaged or poorly imaged retinas, in both small and large-FOV images. The road ahead includes refining the constraint satisfaction scheme with further and more general constraints (including non-geometric ones), exploring alternative constraint satisfaction techniques, and investigating automatic, image-driven tuning of the number of candidates to be returned by individual detectors.

5. REFERENCES

[1] A. Hoover and M. Goldbaum, "Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels," *IEEE Trans. Medical Imaging*, vol. 22, no. 8, pp. 951–8, 2003.

[2] STARE data, "www.parl.clemson.edu/stare/nerve/," .

[3] A Can, H Shen, J N Turner, H L Tanenbaum, and B Roysam, "Rapid automatic tracing and feature extraction from retinal fundus images using direct exploratory algorithms," *IEEE Trans. Information Technology in Biomedicine*, vol. 3, no. 3, pp. 125–138, 1999.

[4] A Pinz, S Bernogger, P Datlinger, and A Kruger, "Mapping the human retina," *IEEE Trans. Medical Imaging*, vol. 17, no. 4, pp. 606–619, 1998.

[5] S Chauduri, S Chatterjee, N Katz, M Nelson, and M Goldbaum, "Automatic detection of the optic nerve in retinal images," in *Proc. IEEE Int.Conf. on Image Processing*, 1989, pp. 1–5.

[6] A Can, C Stewart, B Roysam, and H L Tanenbaum, "A feature-based robust hierachical algorithm for registering pairs of images of the human retina," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 3, 2002.

[7] X Jiang and D Mojon, "Adaptive local thresholding by verification-based multithreshold probing with an application to vessel detection in retinal images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 1, pp. 131–7, 2003.

[8] H Li and O Chutatape, "Automatic location of the optic disk in retinal images," in *Proc. IEEE Int.Conf. on Image Processing*, 2001, vol. 2, pp. 837–840.

[9] F Zana and J-C Klein, "Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation," *IEEE Trans. Image Processing*, vol. 10, no. 7, pp. 1010–1019, 1998.

[10] R Fletcher, *Practical methods of optimization*, Wiley, 1987 (second edition).

[11] W E L Grimson and T Lozano-Perez, "Model-based recognition and localization from sparse range or tactile data," *Int Journ of Robotics Research*, vol. 3, pp. 3–35, 1984.