Elliptical Local Vessel Density: a Fast and Robust Quality Metric for Retinal Images

L. Giancardo, *Student Member, IEEE-EMBS*, M. D. Abràmoff, *Member, IEEE*, E. Chaum, *Member, IEEE*, T. P. Karnowski, *Member, IEEE*, F. Meriaudeau, K. W. Tobin Jr, *Senior Member, IEEE*

***Abstract*— A great effort of the research community is geared towards the creation of an automatic screening system able to promptly detect diabetic retinopathy with the use of fundus cameras. In addition, there are some documented approaches for automatically judging the image quality. We propose a new set of features independent of field of view or resolution to describe the morphology of the patient's vessels. Our initial results suggest that these features can be used to estimate the image quality in a time one order of magnitude shorter than previous techniques.**

I. INTRODUCTION

Diabetic retinopathy is the leading cause of blindness in the Western world. The World Health Organization estimates that 135 million people have diabetes mellitus worldwide and that the number of people with diabetes will increase to 300 million by the year 2025 [1].

New high-quality mydriatic and non-mydriatic cameras have the potential to greatly improve our ability to effectively screen a large population and warn them of a possible disease which may lead to blindness if not treated promptly. In the recent years a substantial effort of the research community went towards the creation of automatic or semi-automatic systems to detect various type of retinopathy. The results obtained are promising but they are strongly dependent on the initial image quality [2].

An algorithm to automatically estimate the quality of a fundus image is essential to build a robust screening system which can be operated not only by expert ophthalmologists, but also by less trained operators. The overall goal is to make these systems suitable for many environments, as it is already the case for the omnipresent automatic blood pressure monitors.

Manuscript received April 7, 2008. This work was supported in part by the National Eye Institute of the National Institutes of Health (RO1 EY017065). This paper was prepared by the OAK RIDGE NATIONAL LABORATORY, Oak Ridge, Tennessee, 37831 6285, operated by UT-BATTELLE, LLC for the U.S. DEPARTMENT OF ENERGY under contract DE AC05 00OR22725.

L. Giancardo is a Master student of the European VIBOT program (e-mail: giancardol@ornl.gov) supervised by F. Meriaudeau.

M. D. Abràmoff is with the Department of Ophthalmology and Visual Sciences, University of Iowa Hospitals and Clinics, Iowa City, IA 52242 USA and also with the Department of Electrical and Computer Engineering, The University of Iowa, Iowa City, IA 55242 USA.

E. Chaum is the Plough Foundation Professor of Retinal Diseases with the University of Tennessee Health Science Center, 930 Madison Avenue, Suite 731, Memphis TN 38163 (e-mail: echaum@utmem.edu)

T.P. Karnowski and K.W. Tobin, Jr are with the Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA (e-mail: karnowskitp@ornl.gov).

However, the measurement of image quality is not a straightforward task, mainly because quality is a subjective concept which varies even between experts, especially for images that are in the middle of the quality scale. Existing algorithms to estimate the image quality are based on the length of visible vessels in the macula region [3], or edges and luminosity with respect to a reference image [4] [5]. Another method uses an unsupervised classifier that employs multi-scale filterbanks responses [6]. The shortcomings of these methods are either the fact that they do not take into account the natural variance encountered in retinal images or that they require a considerable time to produce a result. While algorithms to assess the disease state of the retina need not be as fast (within reason), fast evaluation of quality is essential to the development of an automatic retinopathy scan system. The operator in front of the machine wants to know almost immediately if she is operating the system correctly without having to wait for the relatively long diagnosis process to take place.

II. APPROACH

The approach described in this paper is partially inspired from the work of Niemeijer et al. [6]. Their algorithm, Image Structure Clustering (ISC), gives information on the presence of local image structures at each pixel of the image by clustering the output of a set of multi-scale filters. This technique allows implicit detection of vessels and the optic disc in a completely unsupervised manner. The quality of the fundus image is estimated by employing the global histogram of the ISC features, which is used as the input for a Support Vector Machine classifier similar to the well known "Bag-of-Words" approach [7] [8]. While this technique has the merit of providing a reliable quality estimation without explicitly localising major anatomical structures, it requires a substantial amount of time to classify a new image [1] and it does not take into account the relative position of the features found.

Our approach is based on the hypothesis that a vessel segmentation algorithm's ability to detect the eye vasculature correctly is directly related to the overall quality of an image. Fig 1 shows the vessels detected by our implementation of [9] in images with different quality. It is immediately evident that the low vessel density in the bottom part of the right image is due to illumination issues. Therefore, if we are able

[1] Niemeijer et al. estimated it to be approximately 30 seconds, a figure that is confirmed by our implementation of the algorithm for images with a size of 768x576

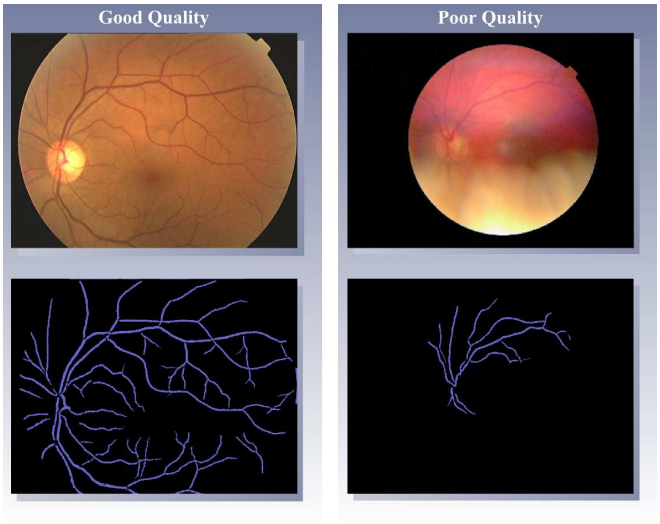

Fig. 1. Comparison of the vessel segmentation in a good and a poor quality fundus image

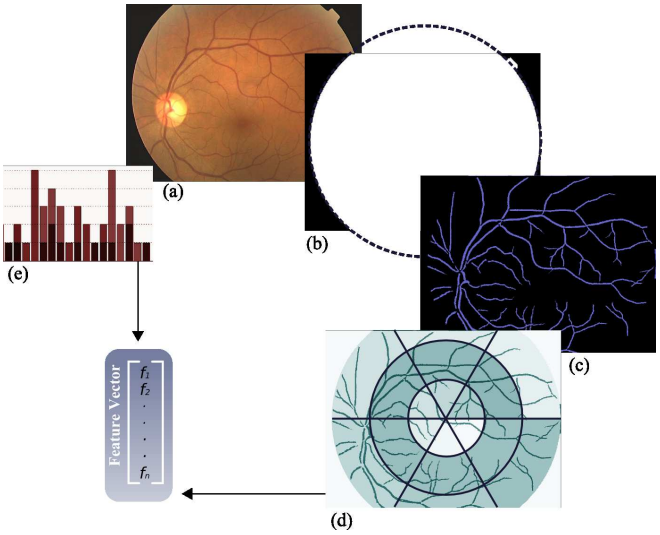

Fig. 2. Visual summary of the feature vector composition. (a) Original Image. (b) Mask detected and fitted ellipse (dashed line). (c) Vessel Segmentation. (d) Local window in polar coordinates. (e) Colour histogram.

to build a fast classifier to estimate the posterior probability that the vessels detected are abnormal, we will be able to judge the quality of the image. The algorithm proposed can be summarised as follows:

$$\begin{array}{c} Mask \quad Localisation \\ \Downarrow \\ Vessel \quad Segmentation \\ \Downarrow \\ Features \quad Vector \quad Composition \\ \Downarrow \\ Classification \end{array}$$

A. Mask Localisation

In this context the mask is a binary image of the same resolution of the fundus image whose positive pixels correspond to the foreground area. The green channel of the RGB fundus image is extracted, and the image is scaled down to 160x120. Four seeds are placed on the four corners of the image and a region growing algorithm is started from these locations. The algorithm measures the mean grey level of the region with $\Delta = 20$ (Δ represents the upper threshold to define common areas in the region growing algorithm). When four regions are segmented, the mask is filled with negative pixels when it belongs to a region and positive otherwise. The process is completed scaling back the image to its original size. The initial and final scaling keep the computational time as small as possible.

B. Vessel Segmentation

The visible eye vasculature is detected using the mathematical morphology method presented by Zana and Klein [9] which proved to be effective in our past work [10]. The algorithm is implemented in C with the OpenCV libraries [11].

C. Feature Vector Composition

The global vessel density could be calculated directly with $\frac{areaVessels}{areaMask}$. Even if this global parameter gives us a strong clue about its importance in a quality estimation context, it is not enough to estimate abnormalities in the vessel segmentation due to the cropping of the image, as seen in the left images in Fig 1.

Our approach tries to normalise any affine homography or image field trimming by fitting an ellipse to the mask. Afterwards, the local vessel densities are extracted employing polar windows which reference the initial ellipse. Each local density is stored in a feature vector, which will be completed by adding colour histogram features. See Fig 2 for a summary.

Ellipse Fitting: The classical technique to fit a geometric primitive like an ellipse to a set of points is the use of iterative methods like the Hough transform [12] or RANSAC [13]. Iterative methods, however, require an unpredictable amount of computational time because the size of the image mask could vary. We employ the non-iterative least squares based algorithm presented by Halir and Flusser [14] which is able to fit more than 100,000 points in a second according to the authors. The points to be fitted by the algorithm are calculated using simple morphological operations on the mask. The complete procedure follows:

$$\begin{array}{c} \alpha \leftarrow erode(maskImage) \\ \gamma \leftarrow maskImage - \alpha \\ fitEllipse(\gamma) \end{array}$$

The erosion is computed with a square structuring element of 5 pixels.

Local Vessel Density: An "adaptable" polar coordinate system (θ, r) is built with the origin coincident with the origin of the ellipse. It is adaptable in the sense that its radius is not constant but it changes according to the shape of the

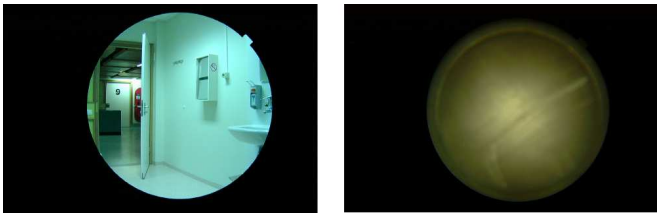

Fig. 3. Examples of Outlier Images

ellipse. The local windows are obtained sampling r every $\frac{r}{3}$ and sampling θ every $\frac{\pi}{3}$, for a total of 18 sections; Fig 2.(d) shows this type of quantisation. The area of the vasculature under each window is computed, normalised with zero mean and unit variance[2] and stored in the feature vector.

Colour Histogram: 5 bins per RGB colour channel of the original image are extracted, normalised with 0 mean and unit variance and added to the feature vector. Now the feature vector will be composed of a total of 33 features. Raw colour histograms are used because they help to distinguish real fundus images from outliers, i.e. images took without patients in front of the camera.

D. Classification

The quality measurement is treated as a classification problem between two classes: "Good Quality" and "Bad Quality". The quality metric used is the posterior probability for the feature vector to belong to the "Good Quality" class. The choice of the appropriate classifier is key for a robust system capable of generalising well, especially if the training set used has a limited number of images. Two families of classifiers are tested: Support Vector Machines (SVM) and K-Nearest Neighbour (KNN). In both cases two external libraries are employed, libSVM [15] for the former and PRtools [16] for the latter.

III. RESULTS

A. Test Data

The dataset employed consists of 84 macula centred fundus images acquired with different resolutions and fields of view. They were divided in 4 classes: "Good", "Fair", "Poor" and "Outliers". The first 3 classes are composed of a subset of the fundus image database provided by M.D. Abramoff [17] [18]. The "Outlier" class consists of images that do not show the patient's eye but which might be erraneously sent to an automatic screening system, examples are shown in Fig. 3. For all the tests 24 images (12 "Good" and 12 "Poor") were extracted from the dataset and used uniquely for the training phase.

B. Global Vessel Density Distribution

The first test performed is an empirical evaluation of our hypothesis about the importance of the vessel density to derive a quality metric. The global histogram of the three image classes is computed using the global vessel density normalised for the mask area using 21 equally spaced bins. Fig. 4 shows the three distributions obtained. Even with this very simple feature it is possible to separate the "Good" from the "Poor" distributions.

[2]The zero mean and unit variance is calculated for each feature across all the training images

TABLE I

CLASSIFIED AS "GOOD QUALITY"

	Actual Class			
Classifier	Good	Fair	Poor	Outlier
SVM (Linear)	100%	83%	0%	11%
SVM (Radial)	100%	91%	0%	11%
KNN (K=1)	100%	66%	0%	11%
KNN (K=8)	100%	83%	0%	66%
ISC, SVM (Radial Kernel)	83%	41%	0%	0%

The first four classifiers employ the ELVD features plus the raw histogram. The percentage represents the number of images correctly identified per class.

TABLE II

QUALITY SCORES

	Actual Class			
Classifier	Good	Fair	Poor	Outlier
SVM (Linear)	0.93 (0.08)	0.6 (0.35)	0 (0)	0.11 (0.2)
SVM (Radial)	0.65 (0.11)	0.43 (0.35)	0 (0)	0.09 (0.15)
KNN (K=1)	0.45 (0.11)	0.26 (0.21)	0 (0)	0.05 (0.08)
KNN (K=8)	0.58 (0.18)	0.45 (0.25)	0 (0)	0.27 (0.21)

The first number represents the average quality score (i.e. The posterior probability for the image to belong to "Good Quality") for all the images tested in a class. The number between brackets is the standard deviation.

C. Quality evaluation

Table I compares the efficiency of four classifiers using our approach with the method presented by Niemeijer et al. [6]. All the classifiers based on Elliptical Local Vessel Density (ELVD) classified inlier images (i.e. real fundus images) with 100% accuracy as opposed to our implementation of the Niemeijer et al. method.

Table II shows a summary of the quality scores obtained for each classifier. In this case, the scores are not directly comparable with the method of [6]. The SVM with linear kernel separates the classes best. The score obtained for the "Fair" class is particularly interesting because the classifier was not explicitly trained with this class. Also, all the classifiers give a score of 0 for the entire "Poor" class.

D. Computational Speed

The computational speed is evaluated on a 2.4 GHz machine with 2 Gb of memory. Every test is run for each image in the dataset and the average time is measured. The Mask localisation phase is implemented in MATLAB and takes an average of 0.49 seconds. Given the iterative nature of the process, we suspect that the speed could be greatly improved when the algorithm is implemented in C++. The Vessel Segmentation phase runs in 2.7 seconds for a single image having a resolution of 756x576 pixels. The algorithm seems to scale well on multi-core machines, as it took roughly the same time to segment two images simultaneously on a Pentium DualCore 2.4. The feature vector composition and

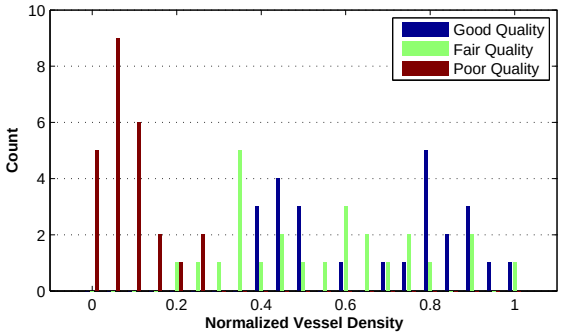

Fig. 4. Vessel Density Distribution

classification phase runs in 0.83 seconds with a SVM (Linear Kernel) classifier. Other classifiers scored a result lower than 1 second. However, if the training set was increased, an exponential computational time increase is expected for the KNN classifier.

IV. CONCLUSIONS AND DISCUSSION

A new set of features for quality assessment macula centred fundus images independent of FOV and scale was presented. The initial results seem to indicate that they can be used by various classifiers to produce a value from 0 to 1 to assign the quality of the image. We compared our approach with an in house implemented version of the state of the art quality classifier described in [6]. The results obtained rule in favour of our technique for inlier images, but not for outliers. The reason is probably due to the initial vessel segmentation technique which mistakes various other structures for vessels. For example, edges of doors (as in Fig 3) are labelled as vessels. Currently, a relatively small image dataset was used with a considerable variation in resolution and FOV, which might explain why the Niemeijer approach did not perform as expected. The next step will be the use of a much larger dataset ground-truthed by different ophthalmologists.

We focused on the computational complexity of our technique. The total running time required to judge the image quality is around 4 seconds. Still, there is ample margin for improvement, especially considering that some of the system components were implemented with an interpreted language (MATLAB).

We would like to note that this work focuses on the quality of the image with respect to the morphology of vascular structure. In future work we would like to numerically prove that the quality classification described can improve the lesion detection algorithms performance if used as a preprocessing step.

Finally we would like to mention that the elliptical local vessel density features seem to have other applications in addition to quality estimation. For example, we divided our dataset in right and left eyes, trained the system with a linear SVM classifier on our feature vectors (without the raw histogram bins) and tested it. Without any optimization, it was possible to correctly identify 92% of right eyes and 100% of left eyes.

REFERENCES

[1] A. F. Amos, D. J. McCarty, and P. Zimmet, "The rising global burden of diabetes and its complications: estimates and projections to the year 2010." *Diabet Med*, vol. 14 Suppl 5, pp. S1–85, 1997.

[2] N. Patton, T. M. Aslam, T. MacGillivray, I. J. Deary, B. Dhillon, R. H. Eikelboom, K. Yogesan, and I. J. Constable, "Retinal image analysis: concepts, applications and potential." *Prog Retin Eye Res*, vol. 25, no. 1, pp. 99–127, Jan 2006. [Online]. Available: http://dx.doi.org/10.1016/j.preteyeres.2005.07.001

[3] A. D. Fleming, S. Philip, K. A. Goatman, J. A. Olson, and P. F. Sharp, "Automated assessment of diabetic retinal image quality based on clarity and field definition." *Invest Ophthalmol Vis Sci*, vol. 47, no. 3, pp. 1120–1125, Mar 2006. [Online]. Available: http://dx.doi.org/10.1167/iovs.05-1155

[4] S. Lee and Y. Wang, "Automatic retinal image quality assessment and enhancement." in *Proceedings of SPIE Image Processing.*, 1999, pp. 1581–1590.

[5] M. Lalonde, L. Gagnon, and M. Boucher, "Automatic visual quality assessment in optical fundus images," *Proceedings of Vision Interface*, 2001.

[6] M. Niemeijer, M. D. Abramoff, and B. van Ginneken, "Image structure clustering for image quality verification of color retina images in diabetic retinopathy screening." *Med Image Anal*, vol. 10, no. 6, pp. 888–898, Dec 2006. [Online]. Available: http://dx.doi.org/10.1016/j.media.2006.09.006

[7] L. Fei-Fei and P. Perona, "A bayesian heirarcical model for learning natural scene categories," in *Proc. CVPR*, 2005.

[8] J. Sivic, B. Russell, A. Efros, A. Zisserman, and W. Freeman, "Discovering object categories in image collections." in *Proc. Int'l Conf. Computer Vision, Beijing*, 2005.

[9] F. Zana and J. C. Klein, "Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation." *IEEE Trans Image Process*, vol. 10, no. 7, pp. 1010–1019, 2001. [Online]. Available: http://dx.doi.org/10.1109/83.931095

[10] K. Tobin, E. Chaum, V. Govindasamy, T. Karnowski, and O. Sezer, "Characterization of the optic disc in retinal imagery using a probabilistic approach," *Progress in biomedical optics and imaging*, 2006.

[11] G. Bradki, A. Kaehler, and V. Pisarevsky, "Learning-based computer vision with intel's open source computer vision library," *Intel Technology Journal*, vol. 09, 2005.

[12] V. F. Leavers, "Shape detection in computer vision using the hough transform," *Springer-Verlag*, 1992.

[13] P. L. Rosin, "Ellipse fitting by accumulating five-point fits," in *Pattern Recognition Letters*, vol. 14, August 1993, pp. 661–699.

[14] R. Halir and J. Flusser, "Numerically stable direct least squares fitting of ellipses," *Department of Software Engineering, Charles University, Czech Republic*, 2000.

[15] C.-C. Chang and C.-J. Lin, *LIBSVM: a library for support vector machines*, 2001. [Online]. Available: http://www.csie.ntu.edu.tw/ cjlin/libsvm

[16] R. P. Duin, D. de Ridder, and P. Juszczak, *PRtools*, 2007. [Online]. Available: http://www.prtools.org/

[17] M. D. Abramoff and M. S. A. Suttorp-Schulten, "Web-based screening for diabetic retinopathy in a primary care population: the eyecheck project." *Telemed J E Health*, vol. 11, no. 6, pp. 668–674, Dec 2005. [Online]. Available: http://dx.doi.org/10.1089/tmj.2005.11.668

[18] M. D. Abramoff, M. Niemeijer, M. S. A. Suttorp-Schulten, M. A. Viergever, S. R. Russell, and B. van Ginneken, "Evaluation of a system for automatic detection of diabetic retinopathy from color fundus photographs in a large population of patients with diabetes." *Diabetes Care*, vol. 31, no. 2, pp. 193–198, Feb 2008. [Online]. Available: http://dx.doi.org/10.2337/dc07-1312