

Abnormality Detection in Retinal Fundus Image Based Splat Feature using Neural Network

V. Sandhiya
M.E. (Embedded Systems)

R. Suguna
AP/ECE
CMS college of Engineering

V. Manoj Kumar
BE (CSE)

M. Saranya
M.E. (Embedded Systems)

Abstract - Retina, an inner layer of the eye consists of rods and cones which provides vision for humans. Haemorrhage is an abnormality in retina and its symptoms are fragile blood vessels, bleeding in blood vessels etc. which affects the vision. In the development of automated screening system, haemorrhage detection in retinal fundus image is an important step. Retinal color images are partitioned in to non-overlapping segments. Each segment is called a splat. The splat contains pixels with similar colour and spatial location. A set of features such as color, contrast, correlation, homogeneity, spatial location, etc are extracted from each splat to describe its characteristics relative to its neighbouring splats. Optimal subsets of splat features are extracted by filter approaches which are followed by wrapper approach. A neural classifier is trained with splat-based expert annotations and evaluated on the publicly available Messidor dataset.

Index terms- haemorrhage detection, splat, and neural classifier

1. INTRODUCTION

The most widely spread and severe eye disease is the diabetic retinopathy. Its symptoms are exudates, drusen, haemorrhage, cotton wool spots and micro aneurysms. This paper concerns only the detection of haemorrhages. There are two different types of haemorrhages, namely small and large haemorrhages. Small haemorrhages are regular in shape and many systems have been developed to detect these lesions, whereas large haemorrhages are irregular in shape and occur infrequently. Their appearances are highly invariable, making a challenge for automated detection.

Retinal haemorrhages are caused by retinal ischemia and are primarily caused by abnormal fragile blood vessels in hypertension, malaria and so on. The evaluation of automated DR detection systems shows that only image containing large haemorrhages provides false negatives of about 50%. Large haemorrhages indicate more severe disease and these improved detection of such lesions will lead to elimination of severe false negatives. Haemorrhage detection primarily fall in to three categories: pixel based approaches, lesion based approaches and image based approaches. Pixel based approaches focuses on the location of haemorrhages on retina. Lesion based approaches use morphological operations to

define candidate lesions and to count them. Image based approaches aimed at detecting eyes with haemorrhages.

Detecting DR lesions are often accomplished by supervised classification, which involves training of classifiers using expert labelled target objects at pixel level. Features are extracted from each pixel and then soft labels are assigned accordingly, indicating the probability of the pixel being one or part of the target object. Abnormal pixels are then combined in to objects. The problems still exists as follows:

- Ideally training samples are intended to be both informative to the classification model and diverse so that information provided by individual samples overlaps as little as possible. But often in a single trained image, there can be a huge number of similar pixel samples.
- It is expensive to acquire expert labelled reference standards for training and evaluation.
- Sensitivity for detection of large haemorrhages has negligible effect on unweighted performance metrics.

The above said problems can be addressed using a higher level entity- the splat, which is a collection of pixels with similar color and spatial location. As haemorrhages consist of blood, they share appearance features with intravascular blood. That makes it difficult to differentiate these from retinal vessels using low level pixel features. On the contrary, by upgrading samples for classification from pixel level to splat level, information is encoded at the splat level, with fewer disturbances from pixel level noise.

The purpose of this study is to present a supervised classification algorithm to detect large, irregular retinal haemorrhages. Reference standard haemorrhage locations were delineated by a retinal specialist (MDA) using splat-based image representation. Supervised classification predicts the likelihood of splats being haemorrhages with the optimal feature subset selected in a two-step feature selection process. From the resulting haemorrhageness map, a haemorrhage index is assigned as the image level output.

2. PRE-PROCESSING

The images that are fed to the systems are acquired at different sites, by different operators using different cameras and camera settings. The first processing step of the system is aimed at making the difference in field-of-view (FOV) size between exams smaller, removing the sharp border between the FOV and the image background and clipping away unused background pixels. A typical image before and after processing is shown in Figure 1. In normal quality images, the FOV may be segmented by thresholding one of the color planes to obtain a binary FOV mask. However, in many cases the use of a fixed threshold will fail due to differences in FOV brightness or the presence of local, underexposed areas in the image. Because there are different shapes of FOV the circularity cannot be used reliably to detect failed segmentations.



Figure 1. Example of an image as it is acquired on site




Figure 2. The same image after pre-processing

The large difference in intensity at the edge of the FOV can be a problem when extracting image features near the border of the FOV. Therefore a “mirroring” operation is performed to remove the intensity gradient at the border. This operation is applied to every pixel outside the FOV.

3. SPLAT SEGMENTATION

Based on the assumption that pixels that are part of the same object or structure sharing similar color, intensity and spatial location, the image is partitioned into non overlapping splats of similar intensity covering the entire image. Splat-based representation is an image re-sampling strategy onto an irregular grid. Background regions with gradual variations in appearance, tends to be consist of fewer

large splats while foreground regions consist of a larger number of smaller splats. At pixel level, the distributions of haemorrhage pixels and non haemorrhage pixels are imbalanced, since haemorrhages usually account for a small fraction of the entire image. Instead of including only a subset of background pixels for training, as many resampling methods do, a splat-based approach maximizes the diversity of training samples by retaining all important Samples.

3.1. Scale-Specific Image Over-Segmentation

Splats are created by over-segmenting images using watershed or toboggan algorithms. Conventional image over segmentation on a regular grid generates so called “super pixels”, a similar concept to “splats.” But super pixels are roughly homogeneous in size and shape, resulting in a lattice pattern. In contrast, a splat-based approach divides images into an irregular grid, depending on the properties of target objects to be detected.

To create splats which preserve desired boundaries precisely, i.e., boundaries separating haemorrhages from retinal background, scale-specific image over-segmentation should occur in two steps. Due to the variability in appearance of haemorrhages, we firstly aggregate gradient magnitudes of the contrast enhanced dark-bright opponency image at a range of scales for localization of contrast boundaries separating blood and retinal background. Next, the maximum of these gradients over scale-of-interest (SOI) is taken in performing watershed segmentation.

Assuming that we establish a scale-space representation of image $I(x,y;k)$ with Gaussian kernel G_k at SOI $k \in k_1, \dots, k_n$, the gradient magnitude $|\nabla I(x,y;k)|$ is computed from its horizontal and vertical derivatives

$$\begin{aligned} |\nabla I(x,y;k)| &= \sqrt{I_x(x,y;k)^2 + I_y(x,y;k)^2} \\ &= \sqrt{\left[\frac{\partial}{\partial x}(G_k * I(x,y))\right]^2 + \left[\frac{\partial}{\partial y}(G_k * I(x,y))\right]^2} \\ &= \sqrt{\left[\frac{\partial G_k}{\partial x} * I(x,y)\right]^2 + \left[\frac{\partial G_k}{\partial y} * I(x,y)\right]^2} \end{aligned} \quad (4.1)$$

$k = k_1, \dots, k_n$

where symbol * represents convolution and $(\partial G_k)/(\partial x)$, $(\partial G_k)/(\partial y)$ are the first order derivatives of Gaussian at scale along the horizontal and vertical direction. The maximum of the gradient magnitude aggregated over the scale band $|\nabla I(x,y)|$ is

$$|\nabla I(x,y)| = \max_i |\nabla I(x,y;k_i)| \quad (4.2)$$

The application of gradient magnitude from a maximum pooling operation $|\nabla I(x,y)|$ across certain scales as the topographic surface in watershed segmentation is important to obtain meaningful splats preserving haemorrhage boundaries precisely. A comparison of it with the original intensity image and gradient images outside SOI as the topographic surface for splat creation is given. Each image in

this figure contains a similar number of splats generated by the same watershed algorithm. The number of splats in each image is set to be within a limit, which is achieved by thresholding the topographic surface iteratively.

3.2. Splat-Based Reference Standard Acquisition

Supervised algorithms require labelled samples by experts, but it is expensive to acquire such data, because substantial time is required to delineate irregular boundaries of haemorrhages. Any misalignment with true boundaries introduces noise at the training stage. Given the limited number of training samples, it has considerable impact on system performances. This problem may be simplified by splat-based image formulation. We compared both a pixel-based approach and a splat-based approach in Figure 3 using "Truth marker"-an iPad app developed to provide a convenient user interface for clinicians to perform reference standard annotation.

For pixel level annotation illustrated in Figure 3(a), we allow two types of annotations by expert: large haemorrhages and small haemorrhages. Large haemorrhages are indicated by a few points along the boundaries (shown as small circles) and then spline fitting is applied to connect those discrete points as enclosed curves shown in cyan. Small haemorrhages are indicated by a single point shown as a green dot. Thus considerable noise is introduced and the time costs of experts are still high. For splat level annotation shown in Figure 3(b), this process is simplified substantially. Experts only perform a single click in a splat to indicate a haemorrhage splat. As splats preserve haemorrhage boundaries, the resulting reference standard is less noisy.




Figure 3. Sample labelling acquired from expert annotation with: (a) pixel-based approach; (b) splat-based approach

To produce an image level reference standard, images with splat-based annotation from the expert, i.e., images containing haemorrhages are given labels of "1" and the rest are given labels of "0" as they contain no haemorrhage splats.

3.3. Edge Effect Removal

Edge effects due to limited field of view (FOV) and vignetting in fundus photographs have to be addressed to suppress irrelevant responses during feature extraction. This effect is visible in Figure 4. It is conventionally performed in two ways. One is to fill the region outside FOV with the mean

color of the region within FOV. The other possibility is to mirror the FOV outside the FOV. If the artifacts were not completely eliminated, they would interfere with features to be identified. This problem can be easily handled with splat-based image representation as is shown in Figure 4. While features are extracted from all of splats, those containing pixels on the circular boundaries of FOV are excluded from further processing. This avoids abrupt intensity changes across splat boundaries and enables the retention of only splats formed by the real content of the image.




Figure 4. Valid splat coverage

4. SPLAT FEATURE EXTRACTION

Given splats with their associated feature vectors and reference standard labels, a classifier can then be trained to detect target objects. In this study, two categories of features are extracted for splat-based haemorrhage detection as follows: 1) splat features aggregated from pixel-based responses; 2) splat wise features (no aggregation is required).

4.1. Pixel-Based Feature Responses

Color within each splat is extracted in RGB color space and dark-bright (db), red-green (rg), and blue-yellow (by) opponency images, which comprises of six color components in splat feature space. To accommodate color variations across the dataset, we normalize each image according to its dominant pixel values at three color channels, which means most frequent pixel values present in the image, are shifted to the origin of RGB color space. No separate rescaling is performed in order to preserve the ratio between color components.

4.2. Aggregation Of Pixel-Based Responses

Similar to the way splats are created so that haemorrhage boundaries are preserved precisely, splat features are more meaningful when response images exhibit high intra-splat similarity and low inter-splat similarity between target classes. To find the optimal strategy to aggregate pixel responses within each splat and associate it with a single feature value, two approaches are used, resulting in four sets of features.

Firstly, the mean and standard deviation (SD) of filtering response within splat are computed. Taking the above DoG responses R_{DOG} , for example,

$$\begin{aligned}
 f_1^\sigma &= \frac{1}{a_p} \sum_{(x,y) \in \Omega_p} [G_\sigma * I(x,y) - G_{\sigma_0} * I(x,y)] \\
 &= \frac{1}{a_p} \sum_{(x,y) \in \Omega_p} R_{DOG}^\sigma(x,y)
 \end{aligned}
 \tag{4.3}$$

$$f_2^\sigma = \left\{ \frac{1}{a_p} \sum_{(x,y) \in \Omega_p} [R_{DOG}^\sigma(x,y) - f_1^\sigma]^2 \right\}^{1/2}
 \tag{4.4}$$

Where Ω_p represents the set of pixels within splat p with area a_p , G_σ represents Gaussian kernel at scale σ and $\sigma_0 = 0.5$.

5. SPLAT FEATURE SELECTION AND CLASSIFICATION

Feature selection, also known as variable selection, attribute selection or variable subset selection, is the process of selecting a subset of relevant features for use in model construction. The central assumption when using a feature selection technique is that the data contains many redundant features or irrelevant features. Redundant features are those which provide no more information than the currently selected features, and irrelevant features provide no useful information in any context. Feature selection techniques are a subset of the more general field of feature extraction. Feature extraction creates new features from functions of the original features, whereas feature selection returns a subset of the features. Feature selection techniques are often used in domains where there are many features and comparatively few samples (or data points). The archetypal case is the use of feature selection in analysing DNA microarrays, where there are many thousands of features, and a few tens to hundreds of samples. Feature selection techniques provide three main benefits when constructing predictive models:

- Improved model interpretability.
- Short training times.
- Enhanced generalisation by reducing the overfitting.

Feature selection is also useful as a part of the data analysis process, as shows which features are important for prediction, and how these features are related. A feature selection algorithm can be seen as the combination of a search technique for proposing new feature subsets, along with an evaluation measure which scores different feature subsets.

5.1. Two-Step Splat Feature Selection

Feature selection reduces the dimensionality of feature space by identifying relevant features and ignoring those irrelevant or redundant ones, which is particularly important to higher separability between classes. There are two major approaches for feature selection: the filter approach and the wrapper approach. The filter approach is fast, enabling their practical use on high dimensional feature spaces. It assesses individual feature separately without considering their interactions. The wrapper approach assesses different combinations of feature subsets tailored to a particular classification algorithm at the cost of longer computation time. To take advantage of both approaches, we

use a two-step feature selection process—a filter approach followed by a wrapper approach.

5.2. Preliminary Feature Selection with a Filter Approach

The goal of preliminary feature selection is to exclude those individual features that are not effective or irrelevant in separating haemorrhage or non haemorrhage splats. It relies on general characteristics of the data to evaluate and select relevant feature subsets without involving any chosen induction algorithms.

The training set is further partitioned into a training subset and a testing subset. Given reference standard labels, splats in the training subset are grouped into haemorrhage splats and non haemorrhage splats. The t-test is applied to each feature of the two groups. The $-t$ -values sorted in ascending order are taken as measures of how effective those features are in predicting the correct labels of splats.

The appropriate number of features to be retained is determined by inspecting how it varies with the misclassification error (MCE) using cross-validation. Classification is carried out using quadratic discriminant analysis (QDA), which performs likelihood ratio test under the assumption of multivariate normal distributions. The percentages of misclassified splats on the training subset and the testing subset are plotted as a function of increasing number of sorted features. Overfitting occurs where the error on the testing subset increases while the error on the training subset decreases. The appropriate number of features is chosen according to the turning point where the smallest MCE on the test set is reached right before overfitting begins to occur.

5.3. Feature Selection with a Wrapper Approach

After preliminary selection, irrelevant features are removed. By taking interactions among features into account, a wrapper approach selects optimal combinations of relevant features with their redundancy minimized. Potential combinations are evaluated depending upon certain classification algorithms. A k-nearest neighbour (kNN) classifier is used for this purpose – the same as what we use for testing in the following sections.

Sorted relevant features identified from the filter approach are applied to sequential forward feature selection (SFS), which attempts to select a feature subset that maximizes area under the receiver operating characteristic (ROC) curve (AUC) of the classification system. The accuracy of splat labels predicted by kNN classifier is assessed using leave-one-out cross-validation.

5.4. K-Nearest Neighbour (KNN) Classification

In k-Nearest Neighbour classification, the training dataset is used to classify each member of a "target" dataset. The structure of the data is that there is a classification variable of interest and a number of additional predictor variables. Generally speaking, the algorithm is as follows:

1. For each row (case) in the target dataset (the set to be classified), locate the k closest members (the k nearest neighbours) of the training dataset. A Euclidean Distance measure is used to calculate how close each member of

the training set is to the target row that is being examined.

2. Examine the k nearest neighbours. Assign this category to the row being examined.
3. Repeat this procedure for the remaining rows (cases) in the target set.

As the computing time goes up as k goes up, but the advantage is that higher values of k provide smoothing that reduces vulnerability to noise in the training data. In practical applications, typically, k is in units or tens rather than in hundreds or thousands.

After feature selection, a trained kNN classifier is set up in a "calibrated" feature space with a set of discriminative features and a set of labelled instances. The kNN classifier assigns soft class labels to query splats based on the labels of their k nearest neighbours in the feature space, i.e., those instances in the training set. When neighbours were labelled as being a haemorrhage splat, the posterior probability that the query splat comes from haemorrhage itself was determined by $p=n/k$. The distance for finding the nearest neighbours is measured with Euclidean metric in the optimized feature space. At the testing stage, the system is fully automatic.

The nearest neighbour rule attempts to estimate the posterior probabilities from labelled training samples. A large value of k is desirable to obtain reliable estimates. But only when all of the nearest neighbours are close enough to the query sample, its a posteriori probability can be approximated by the majority labels of its neighbours. Therefore, a compromise has to be made so that the value of accounts for only a small fraction of the training samples.

6. POST-PROCESSING

Classifier performance is enhanced by the inclusion of a two step post-processing stage: the first step is aimed at filling pixel gaps in detected blood vessels, while the second step is aimed at removing falsely detected isolated vessel pixels.

6.1. Assigning Image Level Haemorrhage Index

The ultimate goal of splat feature classification is to develop a haemorrhage detector, indicating whether or not an image is normal, i.e., free of haemorrhages, or abnormal, i.e., containing one or more haemorrhages. When the a posteriori probability of each splat being haemorrhage is determined, a haemorrhageness map can be created for each testing image. It is then upgraded to a single haemorrhage index as image level decision, which can be fused with results from other lesion detectors consisting of a DR screening system. To eliminate spurious responses in haemorrhageness map, firstly low probability responses are suppressed. A desired separability between haemorrhage splats and the background retina can be reached by setting a threshold where a majority of haemorrhage splats receive higher probabilities than non haemorrhage splats.

Secondly, given a limited number of haemorrhage splats coming out from the first step, those neighbouring ones are merged together to form objects. Objects with small areas are removed because they are more likely to be red lesions or

micro aneurysms, which are supposed to be detected by separate detectors.

Thirdly, splats formed by the fovea, whose locations are detected automatically, are masked out to suppress potential false positives. Because detectors consisting of a screening system attempt to help early detection of sight threatening diseases and prevent their progressions among large populations unaware of any abnormalities for their vision, subjects would have noticed if there are any lesions present at the fovea.

The haemorrhage index assigned at image level is simply calculated as the probability summation of the consequently processed haemorrhageness map.

7. DISCUSSION AND CONCLUSION

In this paper, a splat-based feature classification algorithm with application to large, irregular haemorrhage detection in fundus photographs is presented. Neighbouring pixels with similar intensity are grouped into non overlapping splats. A set of features is extracted from each splat to describe its characteristics. These splats are taken as samples for supervised classification in a selected feature space. Splat-based image representation provides an efficient and natural way to model irregular shaped abnormalities in medical images. Aggregating features within splats improves their robustness and stability, as it is resistant to pixel level noise and intensity bias. Moreover, certain high level texture features are only meaningful when considering regions instead of pixels.

Many of the haemorrhages are connected (continuous) with the retinal vessels. Because many of the false positives in our approach are parts of retinal vessel, an alternative approach would be to mask out all blood vessels using one of the common vessel segmentation methods. However, preliminary studies not presented here show that such an approach, attractive at first consideration, also masked out many of the large haemorrhages we are trying to detect in the first place.

Another potential improvement is use of an active learning approach. As we mentioned earlier, one of the problems for supervised classification is its high cost in acquiring labelled data for training. If we design a classifier that can automatically choose examples with the highest classification uncertainty, i.e., at the decision surface boundary, for manual labelling during the learning process, human experts need to label as little data as possible to achieve the same classification confidence.

REFERENCES

- [1] Abramoff M., Garvin M., and Sonka M. (2010) "Retinal imaging and image analysis", *IEEE Rev. Biomed. Eng.*, vol. 3, pp. 169–208.
- [2] Freeman W.T. and Adelson E.H. (1991) "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 9, pp. 891–906.
- [3] Jitpakdee P., Aimmanee P., and Uyyanonvara B. (2012) "A survey on hemorrhage detection in diabetic retinopathy retinal images", in *Proc. 9th Int. Conf. Elect. Eng./Electron., Comput., Telecommun. Inf. Technol. (ECTI-CON)*, Bangkok, Thailand, 2012, pp. 1–4, vol..
- [4] Kohavi R. and John G. (1997) "Wrappers for feature subset selection", *Artif. Intell.*, vol. 97, no. 1–2, pp. 272–324.
- [5] Romeny B.M.T.H. (2003) "Front-End Vision and Multi-Scale Image Analysis: Multi-Scale Computer Vision Theory and Applications", Written in Mathematica. Berlin, Germany: Springer.

- [6] Sanchez C., Niemeijer M., Abramoff M. And Van Ginneken B. (2010) "Active learning for an efficient training strategy of computer-aided diagnosis systems: Application to diabetic retinopathy screening", *Med. Image Comput. Comput. Assist. Intervent.*, vol. 13, no. 3, pp. 603–610.

IJERT