Segmentation of retinal vessels with a hysteresis binary-classification paradigm

Alexandru P. Condurache*, Alfred Mertins

[a]*Institute for Signal Processing, University of Luebeck, Ratzeburger Allee 160 Building 64 D-23562 Luebeck, Germany*

Abstract

Vessel segmentation in photographies of the retina is needed in a set of computer-supported medical applications related to diagnosis and surgery planning. Considering each pixel in an image as a point in a feature space, segmentation is a binary classification problem where pixels need to be assigned to one of two classes: object and background. We describe a paradigm of hysteresis-classifier design that we apply to the problem of vessel segmentation. Before classification, a multidimensional feature vector is computed for each pixel, such that in the corresponding feature space, vessels and background are more separable than in the original image space. Several classifiers that stem from the hysteresis-classifier design paradigm are tested with this feature space on publicly available databases. These classifiers are very fast and achieve results that are comparable or even superior to known dedicated methods. Hysteresis-based classifiers represent a fast and accurate solution for the retinal-vessel segmentation problem.

Keywords: Retina vessels, image segmentation, binary classification, hysteresis methods, vessel enhancement, feature selection

*Corresponding author

Email addresses: alexandru.condurache@isip.uni-luebeck.de (Alexandru P. Condurache), alfred.mertins@isip.uni-luebeck.de (Alfred Mertins)

1. Introduction

Photographies of the retina showing the vasculature are used, for example, to support medical diagnosis in the case of hypertension, diabetes, and the cardiovascular disease [1], but also for intervention planning [2]. To this end, the retinal vessels need to be segmented to compute measures like vessel area and length, vessel width, abnormal branching, and also to provide a localization of vascular structures. The contrast of vessels in the analyzed images is related to the quantity of blood found therein. Hence, small vessels have a weak contrast. Also differences between vessels and background pixels are localized in the vicinity of vessels. The background is usually inhomogeneous and can be locally similar to the vessels. All these properties lead to poor separability between vessels and background. Conversely, vessels are connected structures, as the blood flows from the large vessels through smaller ones to the capillaries.

The purpose of image segmentation is to separate objects from background, where by objects one understands items of interest in the analyzed images. As a binary pattern classification problem, where each pixel needs to be assigned to either the background or the object class, image segmentation is usually afflicted by class skew, as there are often much more background pixels than object ones. Also, in many cases, there is a strong overlap between object and background in the pixel feature space, usually due to the variations of the object- and background-defining properties. Classification with large class skew is a challenging task, especially for accuracy classifiers that strive to achieve a minimal number of false decisions. Classification in the presence of strong overlap is difficult irrespective of the chosen method, however, good results can be achieved if prior knowledge is used. The segmentation of retinal vessels in photographs of the retina is a particularly challenging problem that has received high interest from the machine-vision community [3]. Most vessel-segmentation algorithms are unsupervised and semi automatic [4], but there are some important applications in which supervised methods are well suited and automatic methods are needed, like, for example, screening for diabetic rethinopathy [5, 6].

Owing to the typical separability problems of vessel images, the first step in extracting the vessels is to enhance them, i.e., to increase their contrast and separability from the background. For this purpose, a large variety of techniques has been proposed including matched filters [7, 8],

Hessian measures [9], the wavelet transform [10, 11], line detectors [12], and morphological image processing [13]. Often these methods include multiscale computations to compensate for the large variety of vessel sizes [14]. Particularly successful have been approaches where various vessel measures, aiming at different vessel properties, have been combined into multidimensional pixel-feature vectors [15, 16, 12, 6] or into new vesselness measures [17, 18].

Vessels can be then segmented in a supervised manner based on a pre-labeled set of examples, for example, using k-nearest neighbor classification [15], hysteresis classification [19], centerline detectors [17], Gaussian mixture model based Bayesian classifiers [11], Support Vector Machines [12], and Neural Networks [6]. Alternatively, unsupervised methods like tracking [10], clustering [16, 20], centerline detection and region growing [21, 14], as well as unsupervised hysteresis classification [22, 23] and multi-thresholding methods [24, 8] can be used. In particular the multi-threshold methods can achieve good results, as they can be specially designed to deal with the large overlap of the vessel and background pixel classes.

In this paper, we describe a new approach to classifier design for binary classification which yields methods well suited for working in a skewed feature space and makes use of prior knowledge to achieve robustness with respect to strong class overlap. We call this approach the *"hysteresis binary-classification paradigm"* and use it to design several classifiers for segmenting vessels in images of the retina. Our paradigm yields fast and accurate classifiers that return a label for each individual pixel. It can be used to generate both supervised and unsupervised methods for either scalar or vectorial inputs. It uses two classifiers: the first one, called the pessimist, works with a practically zero false-positive rate, which with overlapping classes implies a high false-negative rate; the second one, called the optimist, works with a practically zero false-negative rate and a high false-positive rate. Then, using the prior knowledge about the connectivity property of vessels, the pessimist classification can be used to select true vessels from among the optimist classification.

A hysteresis classifier for scalar inputs (i.e., a hysteresis threshold) was also used in [23] in a contribution parallel to ours [22]. However, no consideration was given there to the way the two classifiers should be chosen in relation to the available data (see also Section 2). By design, hysteresis classification yields for scalar inputs a multi-threshold method. Typical multi-threshold

methods are actually local thresholds, where various image regions (varying in size from large portions of the input image to one pixel) get different thresholds when conducting segmentation, as a way to compensate for varying background. The hysteresis threshold, however, always generates just two global thresholds. The final segmentation is then achieved with the help of the connectivity relationship among object pixels. This last connectivity-based step is in a way similar to region growing techniques, where the seeds are generated by the pessimist and the optimist helps to define the logical predicate needed for the growing process. However, the hysteresis methods are significantly faster, as they work with global thresholds. In addition to the algorithm description, we present methods to properly establish the optimal parameter choices for hysteresis classifiers based on training. We obtain results that are comparable or even slightly better, but computed faster, than some of the most powerful state-of-the-art methods specially designed for this problem [8, 24, 15, 11, 21, 12, 18, 6], .

This paper draws on previous work published in conference papers [25, 22, 16, 19] and includes a more complete description and a more thorough evaluation. We use here the vessel-segmentation problem to present a uniform view of the hysteresis binary-classification paradigm. New results and comparisons demonstrate the validity of our approach.

The paper is organized as follows: in Section 2 we describe the hysteresis classification paradigm and its application to vessel segmentation. In Section 3 we describe the feature extraction process that results in the computation of a feature vector for each pixel of an analyzed image. In Section 4 we describe the experiments conducted to test various hysteresis classifiers and compare them to state-of-the-art methods. Finally, in Section 5 we present a critical review of the proposed methods.

2. A hysteresis binary-classification paradigm applied to image segmentation

The concept of hysteresis thresholding for image analysis is not new, being used successfully for edge segmentation by Canny [26] and even before this to construct motion masks [27]. With respect to the design of the hysteresis threshold, Canny only mentions that the high-confidence threshold should be some two times larger than the low-confidence one, as in the edge map the edges are always brighter than the background. Such considerations are characteristic for the way the hysteresis threshold is currently used: the two thresholds usually stay in a fixed, predetermined

relationship to one another, as for example in the work of Niemistö et al. [23]. In this contribution, we introduce a set of rules that permits both the automatic (i.e., supervised) and semi-automatic (i.e., expert-knowledge based and/or unsupervised) design of hysteresis classifiers. The classifiers themselves consist of a high- and a low-confidence classifier for scalar and vectorial inputs and assume that the condition about connectivity of the object-class pixels holds. We use the proposed set of rules to design classifiers for the particular problem of image segmentation and then apply them to the segmentation of vessels.

2.1. Hysteresis classification

In the hysteresis classification paradigm, the high- and low-confidence classifiers are called the pessimist and the optimist, respectively, and they represent the base classifiers. When applying the paradigm to image segmentation, we define connectivity over neighborhood relationships among pixel sites. We introduce a border-separability constraint that is supposed to ensure that the connectivity condition leads exclusively to the selection of object pixels. For image segmentation, our paradigm yields two types of classifiers, the absolute one, taking into consideration only individual pixels, and the relative one, considering also the analyzed image. Finally, for image segmentation we show how to conduct feature extraction using object maps. A flowchart of the hysteresis-classification paradigm is shown in Figure 1.

Base classifiers. If the two classes of a binary classification problem overlap strongly but not completely in some feature space (denoted next by A), then error-free classification is impossible there. If the components of one class do exhibit some type of connectivity in a different feature space (denoted next by B), where there is no overlap (i.e., the classes are disjoint), then the hysteresis paradigm is used to design methods that may achieve error-free classification. Two classifiers working in feature space A (i.e., the pessimist and the optimist), coupled over the connectivity constraint in B, build a *hysteresis classifier*. However, if the connectivity and disjointness in B can be described by some numerical features, then these should be included in A, where a standard classifier can then be used to achieve error-free classification. Thus, we consider the cases in which the connectivity and disjointness in feature space B cannot be easily described numerically.

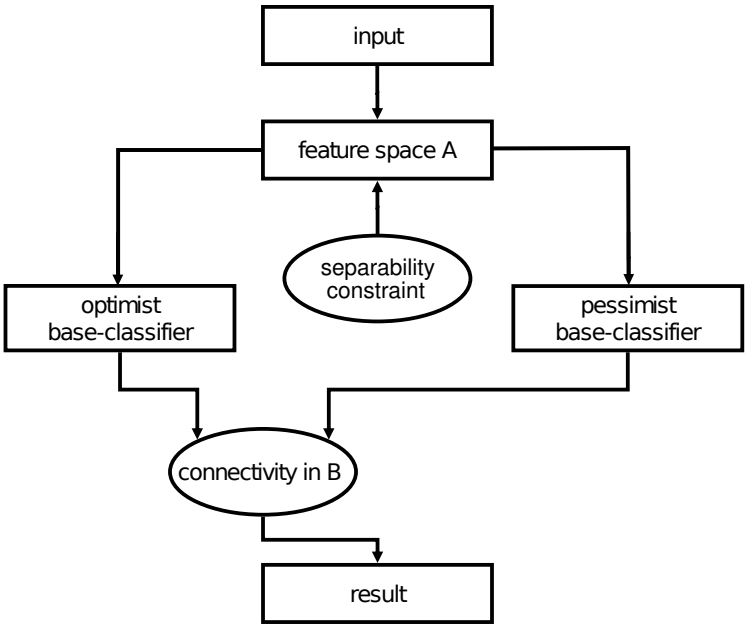

Figure 1: Flowchart of the hysteresis classification paradigm. The border separability constraint ensures that the high-specificity pessimist classification result from feature space A selects, over the connectivity in feature space B, only true object points from the high-sensitivity optimist result.

Hysteresis classification for image segmentation. For the particular case of image segmentation, the features in feature space A are computed from the intensity levels of the analyzed image, such that for each pixel there is a point in the feature space. The feature set B is given by the set of pixel sites $S = \{s_1, s_2, \ldots, s_N\}$ with $N = m \times n$ for an image with m lines and n columns. Each pixel site is represented by a 2D position vector $s_i = [x_i, y_i]^T$ with $i \in \{1 \ldots N\}$. A neighborhood system $\mathcal{N}^k = \{\mathcal{N}_s, s \in S\}$ is introduced as a collection of *pairs* of pixel sites $\mathcal{N}_s$. For a site s_A, a pair is defined as $\mathcal{N}_{s_A} = \{s_A, s_B\}$, such that $s_B \in \mathcal{N}_{s_A} \Leftrightarrow s_A \in \mathcal{N}_{s_B}$ and the distance between them satisfies $d(s_A, s_B) \leq d_{\mathcal{N}}$, with $d(s_A, s_B) = [(x_A - x_B)^2 + (y_A - y_B)^2]^{\frac{1}{2}}$ being the Euclidean distance. If $d_{\mathcal{N}} = 1$, then we have a four-points neighborhood $\mathcal{N}^4$, as there are four pairs for each site. If $d_{\mathcal{N}} = \sqrt{2}$ we have an eight-points neighborhood $\mathcal{N}^8$.

The two base classifiers return two different sets of labels $L = \{l_s, s \in S\}$ with $l_s \in \{0, 1\}$, constituting two different label configurations $L_{optimist}$ and $L_{pessimist}$. With respect to a set of labels, B fulfills a disjointness condition, as each site receives one of two possible and mutually exclusive labels: object ($l_s = 1$) or background ($l_s = 0$). The object connectivity in B is expressed as a proximity relationship, i.e., object points are neighboring other object points according to their

labels and the neighborhood system $\mathcal{N}^k$.

For vessel segmentation, we consider vessels to be connected objects, with each vessel-pixel site being linked over an $\mathcal{N}^8$ neighborhood to another vessel-pixel site. Therefore we segment an image by selecting from $L_{optimist}$ all those vessel-pixel sites connected over a chain of $\mathcal{N}^8$ neighborhoods to a vessel-pixel site from $L_{pessimist}$.

The border-separability constraint. In order to select all object points from A, assuming overlapping classes, the "optimist" will falsely classify many background points as object points. For the hysteresis classification to work without without making any errors, these false objects must be situated in B at a distance larger than the chosen threshold $d_{\mathcal{N}}$ from true objects. In other words, they may not be connected to true objects, and more precisely, they may not be in the neighborhood of true objects. Therefore, the borders of objects have to be sufficiently separable from the surrounding background such that there exists a classifier which yields a segmentation result where true object pixels are separated from false object pixels by $d_{min} > d_{\mathcal{N}}$ in B. We call this property *the border-separability constraint*. In practice, of course, the border-separability constraint will not always hold, and it should be enforced prior to segmentation in the feature-extraction process.

Relative and absolute hysteresis classification. For image segmentation, our paradigm returns two types of hysteresis classifiers: the absolute and the relative classifier respectively. For the absolute hysteresis classifier, we find during training the parameters for two separation surfaces, one for each base classifier. During testing and operation, these surfaces are used to decide for each pixel, irrespective of the image it comes from, whether it is a background or an object pixel. For the relative hysteresis classifier, we find during training two percentile values. These values are used during testing and operation to compute for each analyzed image two individual separating surfaces that are in turn used to compute for each pixel in the respective image whether it belongs to the background or to the object class.

Feature extraction: the object map. In the case of image segmentation, we carry out a chain of processing steps aiming at improving the separability between the two pixel classes, like, for example, improving the contrast and the homogeneity of the pixel-intensity-level representation of

objects and background. For hysteresis segmentation, special care has to be taken during processing to preserve and enhance the objects' borders, in support of the border-separability constraint. The result of the processing chain is an enhanced image called object map, in which each pixel has an intensity value that can be seen as a feature. The generation of the object map will in the following be referred to as a feature extraction process. To avoid undesired bias, the features should be normalized.

2.2. Relative hysteresis classification and percentiles

To design the base classifiers for relative hysteresis classification, we make extensive use of percentiles. The *k*'th percentile is defined as that value of a 1D random variable, which is larger than *k* percent of all other realizations in the available sample. As we are on the real axis, it is self-evident that the two margins of the sample are the maximal and the minimal value. Therefore, the percentile spans the real axis between these two extreme values, defining thresholds to select in uniform steps percentages of the number of realizations in the sample.

For multidimensional feature spaces, we describe here two ways to design the base classifiers: (i) by extending the notion of percentile with the help of a linear classifier and (ii) by applying Linear Discriminant Analysis (LDA) [28] [29] and thus obtaining–in our binary-classification setup–a transformation that maps the original feature space to a line, where we can again work with the usual percentiles.

Percentile-based base classifiers. For the pessimist, the null hypothesis is that the pixel under investigation belongs to the background class. Hence, we impose $P(x_b < t_p) = \alpha$ with x_b being an intensity value in the background class, t_p a threshold, and α the significance. We then have

$$P(x_b < t_p) = \sum_{i=v_{bmin}}^{t_p} \frac{n_{bi}}{N_b} = \alpha, \tag{1}$$

where v_{bmin} represents the minimum intensity value on the histogram of the background intensities, n_{bi} denotes the number of background pixels with intensity value i, and N_b is the total number of background pixels in the image. The value t_p is then the α'th quantile of the histogram of the background's pixel-intensity levels.

The histogram of the image is the discrete approximation of the mixture of vessel and background class-conditional probability density functions (pdfs). Therefore, t_p is also a quantile of the histogram of the image and can be found via

$$P(x < t_p) = \sum_{i=v_{min}}^{t_p} \frac{n_i}{N} = \alpha_{im} \tag{2}$$

where x is a pixel intensity level in the image, v_{min} is the minimum intensity level on the histogram, n_i is the number of pixels with intensity level i, and N denotes the total number of pixels in the image. The threshold t_p is then the α_{im}'th quantile of the histogram of the image, and it should be chosen such that it selects practically *only* vessel pixels.

Similarly, the optimist is computed using the object's class-conditional pdf. This time, we hypothesize that the pixel under consideration is an object pixel. To compute the threshold, we impose again a small significance level β,

$$P(x_o > t_o) = \sum_{i=t_o}^{v_{omax}} \frac{n_{oi}}{N_o} = \beta, \tag{3}$$

where x_o is a pixel-intensity level in the object class, v_{omax} is the maximum intensity level on the histogram of the object's pixel-intensity levels, n_{oi} is the number of object pixels with intensity level i, and N_o is the total number of object pixels in the image. The value t_o is some quantile of the histogram of the object's pixel-intensity levels and it is also a quantile of the histogram of the image. It can be found from

$$P(x < t_o) = \sum_{i=v_{min}}^{t_o} \frac{n_i}{N} = \beta_{im}. \tag{4}$$

The threshold t_o is then the β_{im}'th quantile of the histogram of the image, and it should be chosen such that it selects practically *all* vessel pixels. For the purpose of hysteresis classification we use percentiles (i.e., 100'th quantiles). A histogram with the thresholds t_p and t_o is shown in Figure 2(a), in which the two class-conditional pdfs can be also seen.

The linear-classifier percentile. We introduce the linear-classifier percentile for the 2D case. An extension to more dimensions is straightforward, with the linear-classifier percentile turning from a line into a hyperplane.

To define the linear-classifier percentile, one should first establish the margins of the sample. Optimally, these are along the axis of largest separability. Then, one should also define a way to select percentages of the total number of realizations in the sample. For this second purpose, we need a type of separating surface. A "linear" percentile is obtained when this separating surface is a line. Thus, a linear-classifier percentile is defined by a linear separating surface and by its position on an axis perpendicular to it, i.e., the axis of largest separability. An example is shown in Figure 2(b).

A linear separating surface

$$h_l(\mathbf{x}) = \mathbf{b}^T\mathbf{x} + c = 0 \tag{5}$$

is defined in 2D by the vector of weights $\mathbf{b} = [b_1, b_2]^T$ and the position c. By modifying c, the separating surface is moved over certain distances on the axis defined by $\mathbf{b}$, such that it selects percentages of the available sample in unit steps.

We need to define also a direction on this axis, i.e., an orientation of the separating surface. We define this from the mean of the object class towards the mean of the background class, which is equivalent in the 1D case to considering objects dark. Therefore, the k'th linear-classifier percentile separates k percent of the data from the rest, and most of this data will belong to the object. Then, the pessimist and the optimist are chosen from the set of decision surfaces given by the linear-classifier percentiles from zero to 100. The weight vector $\mathbf{b}$, which is also shown in Figure 2(b), encodes both the axis of largest separability and the orientation of the separating surface.

2.3. Hysteresis classifiers for vessel segmentation

Next we discuss hysteresis thresholds for scalar inputs and hysteresis classifiers for vectorial inputs. A hysteresis threshold works directly on a vessel map. A hysteresis classifier works on a collection of vessel maps such that each pixel is described by a vector (see Section 3). For image segmentation in general, our paradigm yields two types of classifiers, the absolute and the relative hysteresis classifiers that we describe in more detail next.

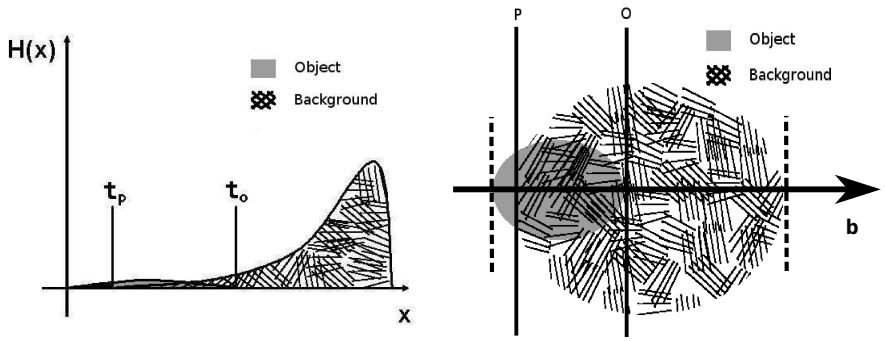

(a) (b)

Figure 2: Schematic representation of an object-map histogram with the pessimist and optimist thresholds, t_p and t_o respectively (a) and schematic representation of a 2D pixel feature-space with the axis of largest separability $\mathbf{b}$ (that ideally goes through the two class means), the margins of the sample (dashed line), the optimist and pessimist classifiers, O and P respectively (b).

2.3.1. Thresholding with hysteresis

When segmenting by thresholding, in a majority of cases a single fixed, global threshold [30] is used for all pixels. There are also adaptive approaches where several thresholds are employed [24], [31], and, in the limit, each pixel obtains its own threshold [32]. The hysteresis paradigm in this case yields an adaptive approach [25] with two thresholds, each corresponding to a base classifier. The threshold can be established automatically or semi-automatically.

2.3.2. The absolute hysteresis classifier

The absolute hysteresis classifier is an LDA-based automatic method. It needs a labeled set of examples from the two classes, viz. object and background. This set is extracted from the pixel feature space.

The pixel feature space. Feature vectors of all pixels from all training images build a pixel-feature vector space. Either the entire vector space, or a sample from it is used to compute the parameters of the hysteresis classifier. These parameters remain constant for all analyzed images.

Inputs: all images in the training set, `number_iterations`, `classifier_type`
$idx_{max} \leftarrow$ `number_iterations`;
$idx \leftarrow 0$;
pessimist(idx)$\leftarrow 0$;
classifier_set = RAW_CLASSIFIER($l = 0$ to 100, `classifier_type`);
$ROC \leftarrow$ BUILD-ROC(classifier_set);
optimist(idx) $\leftarrow \arg\max_{l}$(DISTANCE-TO-ROC-BASELINE(ROC, l));
while $idx \leq idx_{max}$ **do**
  $idx \leftarrow idx + 1$;
  pessimist_classifier_set = RAW_CLASSIFIER($k = 0$ to 100, `classifier_type`);
  $hROC \leftarrow$ BUILD-HYSTERESIS-ROC(pessimist_classifier_set, optimist($idx - 1$));
  pessimist(idx) $\leftarrow \arg\max_{k}$(DISTANCE-TO-ROC-BASELINE($hROC, k$));
  optimist_classifier_set = RAW_CLASSIFIER($m = 0$ to 100,`classifier_type`);
  $hROC \leftarrow$ BUILD-HYSTERESIS-ROC(optimist_classifier_set, pessimist(idx));
  optimist(idx) $\leftarrow \arg\max_{m}$(DISTANCE-TO-ROC-BASELINE($hROC, m$));
  if pessimist(idx) == pessimist($idx - 1$) & optimist(idx) == optimist($idx - 1$)
    exit loop
  end if
end while
return hysteresis-classifier(pessimist,optimist)

Figure 3: Pseudo-code for the iterative training algorithm. Depending on the `classifier_type`, for absolute hysteresis classifiers, the RAW_CLASSIFIER(d,`classifier_type`) is a threshold. For relative classifiers, it is a percentile. RAW_CLASSIFIER(d,`classifier_type`) denotes a classifiers that separates $d\%$ of data from the rest $(100 - d)\%$. By varying d from 0 to 100 we obtain a set of classifiers.

Absolute hysteresis classification and LDA. The absolute hysteresis classifier uses LDA. In [19] a supervised absolute hysteresis classifier is described, where the pessimist and the optimist are two Fisher's classifiers with parameters $(\mathbf{w}, T_p)$ and $(\mathbf{w}, T_o)$ respectively. For our binary classification problem, $\mathbf{w}$ defines an LDA transformation from a multidimensional pixel-feature space to a 1D feature space. $T_p,\ T_o$ represent thresholds in this 1D space. $\mathbf{w}, T_p,$ and T_o represent the parameters of the hysteresis classifier.

During LDA, one looks for a transformation such that in the transformed space, the separability criterion

$$F = \frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2} \tag{6}$$

where $\mu_1,\ \mu_2$ and $\sigma_1,\ \sigma_2$ are the means and variances in the transformed space, is optimized. Using

a labeled input feature space, the transformation weights are

$$\mathbf{w}^T = (\mathbf{m}_1 - \mathbf{m}_2)^T \Big(\frac{n_1}{n} \mathbf{\Sigma}_1 + \frac{n_2}{n} \mathbf{\Sigma}_2 \Big)^{-1} \tag{7}$$

where $\mathbf{m}_1,\ \mathbf{m}_2$ and $\mathbf{\Sigma}_1,\ \mathbf{\Sigma}_2$ are the class-conditional means and covariance matrices, respectively. n_1 and n_2 are the numbers of components in each class and n is the total number of components. $\mathbf{w}$ is computed in the pixel-feature vector space, as described above. T_p and T_o are found during hysteresis training on the pixel-feature vector space.

2.3.3. The relative hysteresis classifier

The relative hysteresis classifier is defined with the help of percentiles. Each image i is considered independently, hence the sample over which the percentiles are computed is given by all pixels in one image at a time. Thus, even if the percentile values obtained during training remain constant for all analyzed images, the parameters of the corresponding separation surfaces change from image to image, thus better adapting to the analyzed data and ultimately providing better results. In the following, we describe two types of relative classifiers: the first uses the linear-classifier percentile [33] and the second is LDA-based [34]. The percentiles are found during hysteresis training.

Linear-classifier percentile-based relative classifier. We use two linear-classifier percentiles as base classifiers. Thus, for an image i we have

$$\begin{aligned} h_p^i &= \mathbf{b}^T \mathbf{x} + c_p^i \\ h_o^i &= \mathbf{b}^T \mathbf{x} + c_o^i \end{aligned} \tag{8}$$

The parameters c_p and c_o represent the positions of the corresponding linear separation surfaces along the axis defined by $\mathbf{b}$ such that they separate certain percentages of the available sample from the rest.

To compute $\mathbf{b}$ we make the assumption that in each image i of the training set the class-conditional pdfs of the object and the background are Gaussian, with equal covariance matrices.

We obtain

$$\mathbf{b}^i = 2\mathbf{\Sigma}^{-1}(\mathbf{m}_b^i - \mathbf{m}_o^i) \tag{9}$$

with $\mathbf{m}_b^i$ being the mean of the background class and $\mathbf{m}_o^i$ denoting the mean of the object class. The vector $\mathbf{b}$ is computed as

$$\mathbf{b} = \frac{1}{M}\sum_{i=1}^{M}\mathbf{b}_i, \tag{10}$$

where M denotes the number of images in the training set. The two percentiles that yield c_p^i and c_o^i for each image are found during hysteresis training based on the images of the training set.

LDA-based relative classifier. The scalar product $\mathbf{b}^T\mathbf{x}$ can also be seen as applying a transformation $\mathbf{b}$ to the data vector $\mathbf{x}$ that maps it to a scalar value. For this purpose we use now the LDA. The base classifiers are now Fisher classifiers, but defined in a relative manner, as the corresponding thresholds T_p^i and T_o^i are computed from percentile values that are set during hysteresis training on the images of the training-set. The vector $\mathbf{w}_i$ is computed for each image i in the training set, and then, $\mathbf{w}$ is computed as the mean over all $\mathbf{w}_i$, similar to the way $\mathbf{b}$ is computed in (10):

$$\mathbf{w} = \frac{1}{M}\sum_{i=1}^{M}\mathbf{w}_i. \tag{11}$$

2.4. Hysteresis training

The purpose of training is to establish the parameters of the base classifiers: for absolute hysteresis classifiers, two thresholds have to be determined, and for relative hysteresis classifiers, two percentile values are needed. If no labeled training set is available, these parameters should be set manually based on some prior knowledge with respect to the segmentation problem. This knowledge can be available, for example, in form of input from an expert or in form of an unlabeled set of training images from which an untrained user can form an opinion for himself. We say we have then a semi-automatic hysteresis classifier. Conversely, we have an automatic classifier and we show how to use the images in the labeled training set to compute the sought parameters. In the vectorial case, we need first to compute $\mathbf{b}$, or $\mathbf{w}$. For absolute hysteresis classifiers we also need to construct the pixel feature space.

Semi-automatic training. Percentiles are related to the percentage of the image surface covered by pixels whose features are situated behind the percentile value. Therefore, to define the hysteresis classifier, one can use prior information about the approximate surface covered by objects in the analyzed images, information which can be translated into percentile values [22].

The percentile defining the pessimist is rather small and is related to the minimal image area which is *certainly* occupied by vessels – this is the black area under the curve in Figure 2(a). Theoretically, this should be the largest percentile selecting *only* vessels. The percentile defining the optimist should be chosen as $\beta_{vm} = 100 - V$, where V defines the largest percentage of the vessel-map surface that can be covered by vessels and is thus related to the minimal image area which is *certainly* occupied by background – this is the dashed area under the curve in Figure 2(a). Theoretically this should be the smallest percentile selecting *all* vessels.

Automatic training. Now, we assume that a set of labeled examples is available for training. Automatic training is thus supervised training. The training algorithm is shown as pseudo code in Figure 3.

Training absolute classifiers. To compute the two thresholds we use the receiver operating characteristic (ROC). The ROC is determined by varying the threshold T of a Fisher classifier $(\mathbf{w}, T)$ between the minimum and the maximum values in the transformed 1D space and computing each time the percentage of true positives or correct classifications (TP) and that of false positives (FP) on the pre-labeled training set. The thresholds T_p and T_o are determined the following way: We start from a standard Fisher classifier used as "optimist", with T corresponding to the ROC-point that is most distant to the baseline linking the ROC-points with FP= 0% and FP= 100% (i.e., the straight diagonal baseline across the ROC graph, corresponding to random performance) and try all "pessimists" by varying T_p between the minimum and the maximum value, building thus a hysteresis-ROC (this is the BUILD-HYSTERESIS-ROC (pessimist_classifier_set, optimist(i)) method in Figure 3), and choose again the one corresponding to the point most distant to the baseline. We repeat this procedure, this time keeping the "pessimist" constant (this is the BUILD-HYSTERESIS-ROC (optimist_classifier_set, pessimist(i)) method in Figure 3), and continue then

for a predetermined number of steps, or until the thresholds do not change anymore.

Training relative classifiers. First, the ROC of a percentile-based decision is used to initialize the optimist as that percentile corresponding to the point which is most distant from the baseline. The percentile-based decision implies running all percentiles from zero to 100, each time segmenting the available images. We then construct the ROC of this decision from the FP and TP rates of each percentile. Then, the ROC of a hysteresis decision is built, using the previously established optimist and all possible pessimist classifiers corresponding to the percentiles from zero to 100. The pessimist corresponding to the point that is most distant from the baseline is selected. The procedure is repeated, this time for the optimist and then in turn for the pessimist and so on for a predetermined number of steps, or until the base classifiers remain unchanged for two consecutive iterations.

3. Feature extraction for retinal-vessel segmentation

The purpose of feature extraction is to compute vessel maps in which the separability of the vessel and object pixel classes is improved, while obeying the border-separability constraint. In the vessel map, which is generated through a sequence of processing steps from a given vessel image (essentially contrast enhancement), each pixel has an intensity value that can be seen as a feature (see also Section 2.1).

Several maps can be used together in a pixel-based multidimensional description of vessels. In this case we can conduct feature selection to find out which vessel maps should be used.

3.1. Vessel maps

In our experiments we have used five vessel maps, described in more detail in [16, 20, 35]:

1. Bothat: the result of a Bothat transform, used to select contrasted structures of a certain size;
2. Hessian single scale: the first eigenvalue of the Hessian matrix, to select elongated structures;
3. Hessian multiscale: the result of the analysis of the eigenvalues of the Hessian matrix extended in a multiscale approach to also reach the small vessels and improve the homogeneity of the vessel class;

4. Band-pass filter: the result of a band-pass filter, to select only vessels based on their size;
5. Laplacian pyramid: the result of a multi-resolution analysis using the Laplace pyramid, to select only fine-detail vessel structures and improve the homogeneity of the vessel pixel class.

In each vessel map, the vessel pixels are characterized by smaller intensities than the background. Each vessel map has been processed to have unit variance and smallest positive mean, such that all pixel intensities are larger than zero. This has been achieved by normalizing first to zero mean and unit variance and then subtracting the smallest value. An example showing the vessel maps computed for an input image is depicted in Figure 4.

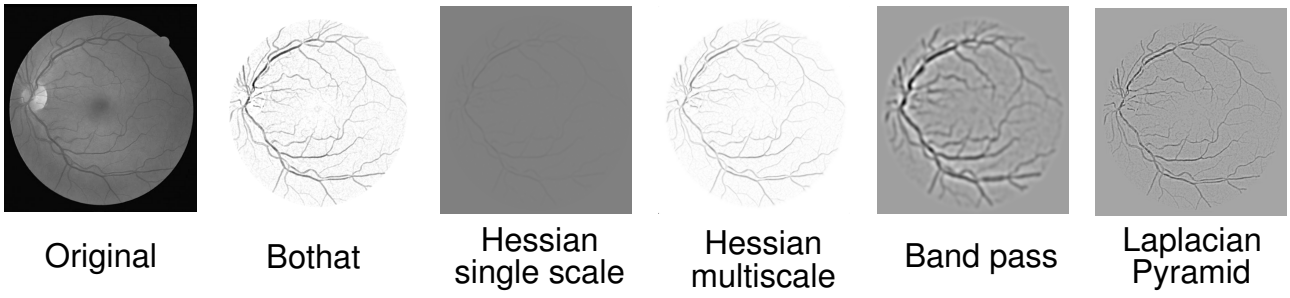

Figure 4: Retinal image and corresponding vessel maps.

3.1.1. Enforcing the separability constraint

Each vessel map depends on a set of parameters. We have chosen these parameters using quality measures defined specifically for hysteresis-based vessel segmentation. These quality measures were designed such as to enforce the separability constraint, i.e., they get larger values for vessel maps where the vessels can be better separated from their immediate vicinity. An example of such a measure is the background-less partial area under the ROC (background-less pAROC). To compute this quality measure for a vessel map, we need a labeled ground-truth image. First we define a ground-truth mask by morphologically dilating the ground-truth image, thus being able to select only the vessels and their immediate vicinity. Next, the vessel map is segmented by a set of thresholds corresponding to the intensity-level percentiles from zero to 100, and a modified ROC (mROC) curve is built. The false-positive rates of the mROC are computed only from the region selected by the ground-truth mask. The quality measure is given by the pAROC, computed using

a 2% bound on the false-positives rate. We chose this pAROC because for "good" vessel maps, we expect the true-positives rate to increase more rapidly at small false-positives rates in the direct vicinity of a vessel. More details and a discussion over various such measures can be found in [35].

3.2. Pixel-based multidimensional description of vessels

A multidimensional feature space is obtained by combining the results of several different vessel-enhancement methods. For each pixel, we build a feature vector by ordering its scalar features (i.e., intensities) in each object map into a vector [16, 19].

Within each object map we seek to increase the separability between the background and object pixel-classes and, if possible, to enforce also the border-separability constraint. Our strategy is in this case to combine the results of several different enhancement methods, in the hope that together they constitute a more separable representation of objects and background than any of them taken alone. It is believed that a multidimensional pixel feature space is better than a single vessel map, as it includes more information about vessels, acquired from different perspectives. A schematic representation of the way a pixel-based feature vector space is computed is shown in Figure 5.

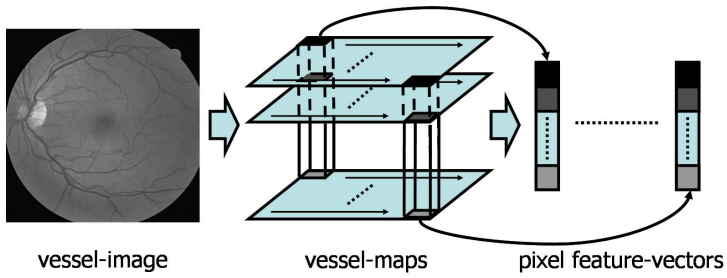

Figure 5: Schematic representation of the method to achieve a pixel-based multidimensional description of vessel and background.

3.2.1. Feature selection

The ROC of various decision rules can be used to characterize the feature space. Clearly, the larger the area under the ROC (AROC), the more separable the feature space is. In the limit the two classes are linearly separable when AROC is one, i.e 100% TP for 0% FP. If several features have been computed, then those which build the best feature space will also yield the largest AROC.

In the case of hysteresis classification we are interested also in the derivative of the ROC curve, especially in the region where the FP rate is small. A hysteresis classifier trained on a ROC curve with a large integral, but a comparatively mild increase over the region with small FP rate, will yield rather poor results, because the "pessimist" classifier will select too few true vessels. Thus, it is better to consider only a pAROC. In our experiments we have computed only the AROC bounded by a 30% FP rate.

Several strategies can then be followed. For example, a full search strategy where all possible combinations of features are investigated, or a sequential search strategy [36], where first the single best feature is selected, then the best combination between that feature and another one and so on, until the optimal feature set is found.

4. Results and discussion

We have used two databases that are publicly available: the DRIVE database [15] and the STARE database [8]. The DRIVE database contains 40 images, divided into a training and a test set. Each of these two sets contains 20 images. For the test images there are two sets of hand-labeled ground-truth images, marked as first and second observer, respectively, as they were generated by different groups of persons. For the training images, there is just one set of ground-truth images, marked first observer. We have used the first-observer set as ground truth in our experiments.

The STARE database contains 20 images. As there is no training and test set, the segmentation performance on this database is computed by means of the leave-one-out method. The STARE database contains two sets of hand-labeled ground-truth images, again marked as first and second observer respectively. We have used the first-observer set as ground truth in our experiments.

Using the training set of the DRIVE database, the parameters of the enhancement methods used to generate the vessel maps were computed such as to optimize a quality criterion for vessel maps [35]. This quality criterion considers both the border-separability constraint and the separability between the two pixel classes. The same parameters were used for the STARE database as well.

For feature selection, in our experiments we use the pixel-feature space built from the training images of the DRIVE dataset, and compute the ROC as in the case of training absolute hysteresis

classifiers: the ROC is determined by varying a threshold between the minimum and the maximum values in the transformed 1D space and computing each time the percentages of TP and FP on the pre-labeled training set. Having only five features to choose from, we have used here a full-search strategy. After feature selection, only the Hessian multiscale, the Hessian single scale and Laplacian-pyramid based vessel maps remained. This type of feature vector was used on the STARE images as well.

The performance of the classifiers was measured by the AROC. The corresponding ROC is computed by fixing the pessimist and modifying the optimist such that it assigns to the vessel class between 0% and 100% of the available test samples. We have also computed accuracy, sensitivity and specificity. The *accuracy* is defined as

$$Ac = \frac{N_{TP} + N_{TN}}{N_{TP} + N_{FN} + N_{TN} + N_{FP}} \tag{12}$$

with N_{TP} being the number of true positives, N_{TN} the number of true negatives, N_{FP} the number of false positives, and N_{FN} the number of false negatives. The *sensitivity* is defined as

$$S_e = \frac{N_{TP}}{N_{TP} + N_{FN}} \tag{13}$$

and the *specificity* as

$$S_p = \frac{N_{TN}}{N_{TN} + N_{FP}} \tag{14}$$

The performance measures rely on the manual ground truth. For image segmentation in general [37], and for vessel segmentation in particular, such a ground truth is difficult to compute. While conducting hand-labeling experiments in our group, with the help of students, image processing experts, and physicians, we have noticed that the human observer usually tends to ignore very thin vessels and to enlarge thick vessels. Therefore, some false positives may still be true vessels, and some false negatives may actually be background. However, such problems are at least partially alleviated by using publicly available databases to conduct comparisons among vessel-segmentation methods.

Table 1 and Table 2 contain the results for the two databases. All results are average values

Classifier type	AROC	A_c	S_e	S_p
Hysteresis absolute	0.9642	0.9484	0.9053	0.9517
Hysteresis relative lin.-prc.	0.9713	0.9509	0.9086	0.9580
Hysteresis relative LDA	**0.9726**	0.9516	**0.9094**	0.9591
Soares et al.(2006) [11]	0.9614	0.9466	-	-
Staal et al.(2004) [15]	0.9520	0.9441	0.7194	0.9773
Jiang and Mojon(2003) [24] (from [18])	0.9327	0.8911	-	-
Mendonca and Campilho(2006) [21]	-	0.9463	0.7315	0.9781
Ricci and Perfetti(2007) [12]	0.9633	**0.9595**	-	-
Lam et al.(2010) [18]	0.9614	0.9472	-	-
Marin et al.(2011) [6]	0.9588	0.9452	0.7067	**0.9801**
Second observer	-	0.9473	0.7761	0.9725

Table 1: Results achieved on the DRIVE database by different segmentation methods. The best values are written in bold characters.

over all test images in the respective database. Some classification examples are shown in Figure 6. In Table 1 we show also results obtained by some state-of-the-art vessel-segmentation methods on the DRIVE database, as reported in the respective articles. The sensitivity and specificity entries for the other methods (except for [6]) have been taken from the work of Mendonca and Campilho [21]. The specificity is not directly mentioned there, but instead the false positives fraction (FPF), defined as the fraction of pixels erroneously classified as vessel points that we interpret as $FPF = 1 - S_p$.

In the following, we discuss various hysteresis classifiers. The relative LDA-based hysteresis classifier shows the best performance on the test images. On the DRIVE database, with respect to sensitivity at a significance level $\alpha = 0.02$ and with respect to specificity at a significance level $\alpha = 0.005$, it is significantly better than the linear-percentile based relative classifier. The observations with respect to the AROC and the accuracy remain valid on the STARE database as well. However on this database, the LDA-based classifier is more specific and less sensitive than the linear-percentile based classifier. This means that it will yield a smaller number of correct classifications but also less false positives. The relative linear-percentile classifier has a slightly more pronounced tendency to oversegmentation in comparison to the LDA-based relative classifier. We interpret this as an indication that the explicit Gaussianity assumption made in the case of the former classifier is less likely to hold.

Classifier type	AROC	A_c	S_e	S_p
Hysteresis absolute	0.9650	0.9569	0.8879	0.9652
Hysteresis relative lin.-prc.	0.9757	0.9574	*0.8907*	0.9654
Hysteresis relative LDA	**0.9791**	0.9595	0.8902	0.9673
Soares et al.(2006) [11]	0.9671	0.9480	-	-
Staal et al.(2004) [15]	0.9614	0.9516	0.6970	0.9810
Jiang and Mojon(2003) [24] (from [18])	0.9298	0.9009	-	-
Mendonca and Campilho(2006) [21]	-	0.9479	0.7123	0.9758
Ricci and Perfetti(2007) [12]	0.9680	**0.9646**	-	-
Lam et al.(2010) [18]	0.9739	0.9567	-	-
Marin et al.(2011) [6]	0.9767	0.9526	0.6944	**0.9819**
Second observer	-	0.9351	**0.8949**	0.9390

Table 2: Results achieved on the STARE database by different segmentation methods. The best values are written in bold characters. In the case of S_e, the value given in italic characters represents the best value achieved by automatic methods.

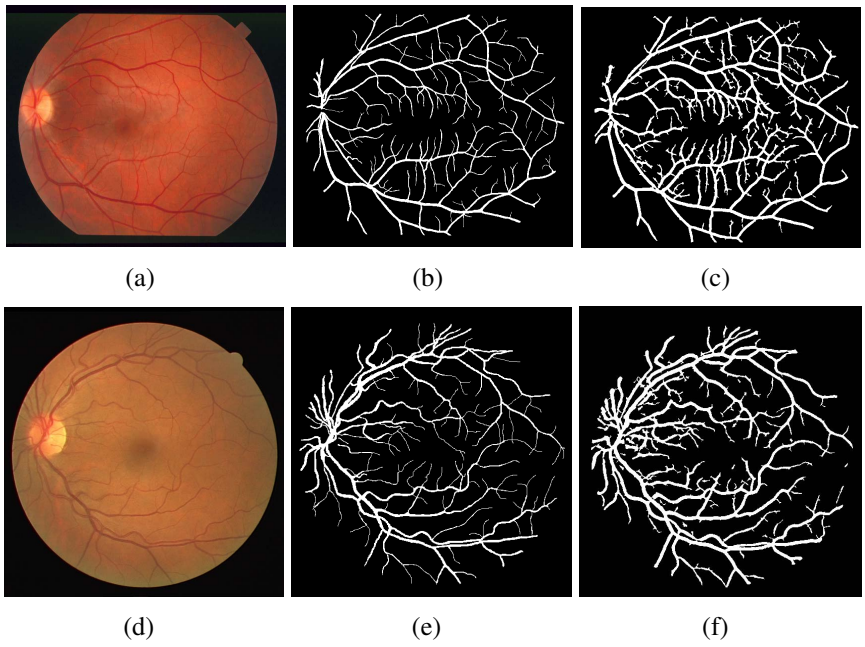

Figure 6: Image of the STARE database (a), the corresponding ground truth (b) and segmentation result achieved by the relative hysteresis classifier (c). Image from the DRIVE database (d), ground truth (e) and result of segmentation by the relative hysteresis classifier (f).

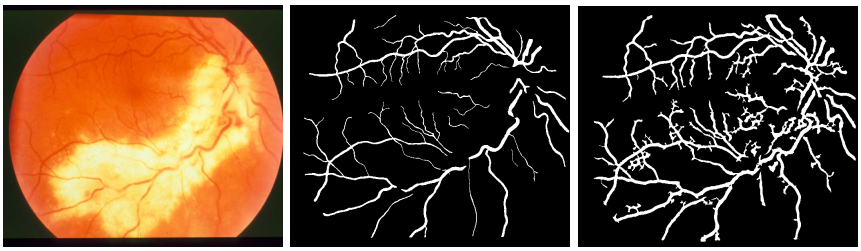

(a) (b) (c)

Figure 7: Result on a pathological image: input image (a), manual ground truth (b) and segmentation result computed with the relative linear-percentile based hysteresis classifier (c).

As the results in Table 1 show, the hysteresis methods have the best sensitivity from among the automatic segmentation methods, but they tend to trade this in for specificity, thus being more likely to return slightly oversegmented results. Nevertheless, they show a very good accuracy, which means that the tradeoff is worth it. The supervised hysteresis classifiers perform better than other vessel-segmentation methods with respect to the AROC. Only the SVM-based method by Ricci et al. [12] has an improved accuracy, but a smaller AROC. This shows that our training algorithm is not yet optimal–in our experiments we have stopped the training after a predetermined number of steps. However, our strategy based on computing pixel-features from vessel maps is most promising, because theoretically, the achieved results could be further improved by improving the separability in the feature space, i.e., designing better vessel maps.

The STARE database includes relatively many images showing pathological retinas. Our vessel-map based hysteresis approach is not designed to handle such images. Nevertheless, it is relatively successful in some cases. Such an example is shown in Figure 7. The strong background variation caused by the large bright spot is reduced during feature extraction to a level that does not afflict the segmentation. Large bright background structures can be eliminated in vessel maps computed for example with the help of morphological image processing algorithms.

On a machine with a Core 2 Duo E6700 processor and 2 GB memory under MATLAB, the training time for a relative LDA-based classifier was one and a half hours on the DRIVE data set. The time needed to reach a result by the same classifier in the test and operation phase is 0.8 seconds per image from the DRIVE database. The time needed to compute the pixel-feature vector set for the same image is six seconds. Therefore a new image is segmented every 6.8

	Training time (hours)	Time (seconds)	PC configuration	Software
relative LDA	**1.5**	**6.8**	Core 2 Duo, 2.66GHz, 2GB RAM	MATLAB
Soares et al. [11]	9	180	Athlon XP, 2.17GHz, 1GB RAM	MATLAB
Staal et al. [15]	–	900	P3, 1GHz, 1GB RAM	–
Jiang and Mojon [24]	–	∼19	P3, 600MHz	C
Mendonca and Campilho [21]	–	150	P4, 3.2GHz, ∼1GB RAM	MATLAB
Lam et al. [18]	–	780	Core 2 Duo, 1.83GHz, 2GB RAM	MATLAB
Marin et al. [6]	–	93	Core 2 Duo, 2.13GHz, 2GB RAM	–

Table 3: Running times for different vessel segmentation methods on the DRIVE database. The best values are written in bold characters.

seconds. The absolute hysteresis classifier is a few milliseconds faster. Hysteresis methods yield classifiers that are significantly faster than other state-of-the-art methods. An overview of running times for various vessel segmentation methods is given in Table 3.

Supervised vessel segmentation can be achieved also by a general-purpose classifier like an SVM working on the pixel-feature space described in this contribution. However, the strong overlap has various negative consequences. First, there will be a rather large number of support vectors, therefore, it should take a relatively long time to segment an image. There is also the danger of overfit, particularly if the size of the sample used for training is too small. Conversely, if the entire available training database is used (i.e., all pixels from all images in the training set), the time needed for training is very long, which makes designing the SVM a tedious job. Ricci et al. [12] successfully use a linear SVM to conduct vessel segmentation (on a different feature space). For training they used only 20000 pixels randomly chosen from all images in the training set. There

is no mention of the time needed to segment one image in their paper, except for a comment that the method should be very fast, being based on a linear SVM. Marin et al. [6] propose neural networks (NNs) to segment vessels on a newly for this purpose designed feature space. Their feature space bares resemblance to ours, as it assigns to each pixel a feature vector consisting of intensity-level-based features, but also of moment-invariants-based features. For training they use a set of 30000 manually selected pixels from the DRIVE database, fairly divided between vessels and background.

5. Conclusions and summary

Hysteresis segmentation can successfully segment objects of inhomogeneous intensity-level representation found on an inhomogeneous background, as long as there is a slight difference between object and background at a local level around the object's borders, and the supports of the two classes in the pixel feature space do not overlap completely. These conditions should be enforced during feature extraction also. Thus, for vessel segmentation, the feature-extraction process was designed to improve the contrast of vessels, particularly at their borders, such that the vessels do not spill into the background when segmented by the optimist.

We do not treat the case of pathological retinal images separately. In general the performance of the methods described in this contribution on pathological images is worse than on normal images. However, the hysteresis paradigm can be used to handle such problems as well. In such a case, special vessel maps need to be devised to improve vessel separability on images with specific pathologies. This represents an indication of the versatility of our vessel-map-based feature-extraction process. The feature extraction is aimed at ensuring the border separability constraint and improving the overall separability.

For retinal vessel segmentation, the hysteresis methods tend to trade specificity for sensitivity. However the good accuracy obtained shows that it was a good trade and that the danger of oversegmentation is small. Furthermore, these considerations need to be weighted by the fact that manually labeled gold-standard images used to compute these quality measures may themselves by afflicted by errors.

Requiring the pessimist to have a zero false positives rate is practically equivalent, in the case of a linear classifier, to selecting that region in the feature space A that intersects with the support of the object class and it can be separated from the support of the background class by a linear hypersurface. If only nonlinear separation is possible, then a nonlinear classifier is needed as pessimist. We intend to investigate nonlinear hysteresis classification in the near future.

Sometimes, the set of separating surfaces computed by percentiles is too sparse such that the optimist and the pessimist can not be defined correctly. For the pessimist, this means that there is no percentile-based separating surface to select only object points, although such a separating surface exists. For the optimist, this means that there is no separating surface to fulfill the border-separability constraint, although such a separating surface exists. Then, instead of using percentiles, one should use other quantiles (larger than 100) which yield a sufficiently dense set of separating surfaces.

We have introduced a new paradigm for binary classification including supervised and unsupervised methods, which is well suited to solve problems afflicted by large class skew and strong overlap by making use of available prior knowledge. Our paradigm is rooted in the field of image segmentation and analysis, and we have demonstrated it successfully for the segmentation of retinal vessels. In this case we have a large class skew, because there are a lot more background pixels than vessel pixels and there is a strong overlap as the vessels and the background are inhomogeneous. Our paradigm successfully uses the prior knowledge that the vessels are connected structures. Even though it has been demonstrated here only for vessel segmentation, we believe the hysteresis paradigm is a more general method, readily applicable to other types of segmentation and binary-classification problems as well.

References

[1] Kansky, J.. Clinical ophthalmology: A systematic approach. London, U.K.: Butterworth-Heinemann; 1989.

[2] Shen, H., Roysam, B., Stewart, C.V., Turner, J.N., Tanenbaum, H.L.. Optimal schedul-

ing of tracing computations for real-time vascular landmark extraction from retinal fundus images. IEEE Transactions on Information Technology in Biomedicine 2001;5(1):77–91.

[3] Kirbas, C., Quek, F.K.H.. A review of vessel extraction techniques and algorithms. ACM Computing Surveys 2004;36(2):81–121.

[4] Pham, D., Xu, C., Prince, J.. Survey of current methods in medical image segmentation. Anual Review of Biomedical Engineering 2000;2:315–337.

[5] Klonoff, D.C., Schwartz, D.M.. An economic analysis of interventions for diabetes. Diabetes Care 2000;23(3):390–404.

[6] Marin, D., Aquino, A., Gegundez-Arias, M., Bravo, J.. A new supervised method for blood vessel segmentation in retinal images by using gray-level and moment invariants-based features. IEEE Transactions on Medical Imaging 2011;30(1):146–158.

[7] Poli, R., Valli, G.. An algorithm for real time vessel enhancement and detection. Computer Methods and Programs in Biomedicine 1997;52(1):1–22.

[8] Hoover, A., Kouznetzova, V., Goldbaum, M.. Locating blood vessels in retinal images by picewise threshold probing of a matched filter response. IEEE Transactions on Medical Imaging 2000;19(3):203–210.

[9] Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.. Multiscale vessel enhancement filtering. Lecture Notes in Computer Science 1998;1496:130–137.

[10] Chen, Z., Molloi, S.. Multiresolution vessel tracking in angiographic images using valley courses. Optical Engineering 2003;42:1673–1682.

[11] Soares, J.V.B., Leandro, J.J.G., Jr., R.M.C., Jelinek, H.F., Cree, M.J.. Retinal vessel segmentation using the 2-D Gabor wavelet and supervised classification. IEEE Transactions on Medical Imaging 2006;25(9):1214–1222.

[12] Ricci, E., Perfetti, R.. Retinal blood vessel segmentation using line operators and support vector classification. IEEE Transactions Medical Imaging 2007;26(10):1357–1365.

[13] Fraza, M., Barmana, S., Remagninoa, P., Hoppea, A., Basitb, A., Uyyanonvarac, B., et al. An approach to localize the retinal blood vessels using bit planes and centerline detection. Computer Methods and Programs in Biomedicine 2011;.

[14] Martinez-Perez, M.E., Hughes, A.D., Thom, S.A., Bharath, A.A., Parker, K.H.. Segmentation of blood vessels from red-free and fluorescein retinal images. Medical Image Analysis 2007;11(1):47 – 61.

[15] Staal, J., Abramoff, M.D., Niemeijer, M., Viergever, M.A., van Ginneken, B.. Ridge-based vessel segmentation in color images of the retina. IEEE Transactions on Medical Imaging 2004;23(4):501–509.

[16] Condurache, A.P., Aach, T., Grzybowski, S., Machens, H.G.. Vessel segmentation and analysis in laboratory skin transplant micro-angiogram. In: Proceedings of CBMS-2005. Dublin, Ireland: IEEE; 2005, p. 21–26.

[17] Sofka, M., Stewart, C.. Retinal vessel centerline extraction using multiscale matched filters, confidence and edge measures. IEEE Transactions Medical Imaging 2006;25(12):1531–1546.

[18] Lam, B.S.Y., Gao, Y., Liew, A.W.C.. General retinal vessel segmentation using regularization-based multiconcavity modeling. IEEE Transactions Medical Imaging 2010;29(7):1369–1381.

[19] Condurache, A.P., Aach, T.. Vessel segmentation in 2D-projection images using a supervised linear hysteresis classifier. In: Proceedings of ICPR-2006; vol. 1. Hong Kong, China: IEEE; 2006, p. 239–243.

[20] Condurache, A.P., Aach, T., Grzybowski, S., Machens, H.G.. Imaging and analysis of angiogenesis for skin transplantation by microangiography. In: Proceedings of ICIP-2005. Genoa, Italy: IEEE; 2005, p. II/1250 – II/1253.

[21] Mendonca, A.M., Campilho, A.. Segmentation of retinal blood vessels by combining the detection of centerlines and morphological reconstruction. IEEE Transactions on Medical Imaging 2006;25(9):1200–1213.

[22] Condurache, A.P., Aach, T.. Vessel segmentation in angiograms using hysteresis thresholding. In: Proceedings of MVA-2005. Tsukuba, Japan; 2005, p. 269–272.

[23] Niemistö, A., Dunmire, V., Yli-Harja, O., Zahng, W., Shmulevich, I.. Robust quantification of in vitro angiogenesis through image analysis. IEEE Transactions on Medical Imaging 2005;24(4):549–553.

[24] Jiang, X., Mojon, D.. Adaptive local thresholding by verification based multithreshold probing with application to vessel detection in retinal images. IEEE Transactions on Pattern Anaysis and Machine Intelligence 2003;25(1):131–137.

[25] Condurache, A.P., Aach, T., Grzybowski, S., Machens, H.G.. Vessel segmentation for angiographic enhancement and analysis. In: Proceedings of BVM-2005. Heidelberg, Germany: Springer; 2005, p. 173–177.

[26] Canny, J.. A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 1986;8(6):679–698.

[27] Connor, D.J., Haskell, B.G., Mounts, F.W.. A frame-to-frame picturephone coder for signals containing differential quantizing noise. The Bell System Technical Journal 1973;52(1):35–51.

[28] Fukunaga, K.. Introduction to statistical pattern recognition. Academic Press; 1990.

[29] Loog, M., Duin, M.P.W.. Liniar dimensionality reduction via a hetereoscedastic extension of lda: the chernoff criterion. IEEE Transactions on Pattern Anaysis and Machine Intelligence 2004;6(26):732–739.

[30] Sezgin, M., Sankur, B.. Survey over image thresholding techniques and quantitative performance evaluation. Journal of Electronic Imaging 2004;13(1):146–165.

[31] Aach, T., Condurache, A.P.. Transformation of adaptive thresholds by significance invariance for change detection. In: Proceedings of SSP-2005. Bordeaux, France: IEEE; 2005,.

[32] Aach, T., Kaup, A.. Bayesian algorithms for adaptive change detection in image sequences using Markov random fields. Elsevier Signal Processing: Image Communication 1995;7:147–160.

[33] Condurache, A.P., Mertins, A., Aach, T.. Supervised, hysteresis-based segmentation of retinal images using the linear-classifier percentile. In: Medizinische Bildverarbeitung und Mustererkennung (GI Jahrestagung); vol. P154 of *Lecture Notes in Informatics*. Lübeck: GI; 2009, p. 1285–1293.

[34] Condurache, A.P., Müller, F., Mertins, A.. An LDA-based relative hysteresis classifier with application to segmentation of retinal vessels. In: Proceedings of ICPR-2010. Istanbul, Turkey: IEEE; 2010, p. 4202–4205.

[35] Condurache, A.P.. Cardiovascular biomedical image analysis: methods and applications. Waabs: GCA; 2008.

[36] Pudil, P., Novovicova, J., Kittler, J.. Floating search methods in feature selection. Pattern Recognition Letters 1994;15(11):1119–1125.

[37] Bouix, S., Fernandez, M.M., Ungar, L., Nakamura, M., Koo, M., McCarley, R., et al. On evaluating brain tissue classifiers without a ground truth. NeuroImage 2007;36(4):1207–1224.