

Computational Methods and Apparatus for Meibography

Field of the invention

The present invention relates to computational methods and apparatus for processing images of the meibomian glands to derive information characterizing abnormalities in the glands, indicative of medical conditions.

Background of the invention

The meibomian glands are sebaceous glands at the rim of the eyelids inside the tarsal plate, responsible for the supply of meibum, an oily substance that prevents evaporation of the eye's tear film. Meibum is a lipid which prevents tear spillage onto the cheek, trapping tears between the oiled edge and the eyeball, and makes the closed lids airtight. It further covers the tear surface, and prevents water in the tears from evaporating too quickly. Dysfunctional meibomian glands can cause dry eyes (since without this lipid, water in the eye evaporates too quickly) or blepharitis; and other medical conditions.

It is known to capture infra-red (IR) images of ocular surface to analyse the morphological structures of the meibomian glands. For an healthy eye, the glands have similar features in terms of spatial width, in-plane elongation, length, etc. On the other hand, for an unhealthy eye, the imaged glands show irregularities. Thus, it is important to detect each gland and extract features such as orientation, width, length, curvature, etc., for the purpose of automated dry eye diagnosis and risk assessment.

Figs. 1(a)-(c) are three sample IR images of an ocular surface. Fig. 1(a) has been graded manually by an expert as "healthy", Fig. 1(b) as "intermediate" and Fig. 1(c) as "unhealthy". In this document such images are respectively referred to as "healthy images", "intermediate images" and "unhealthy images". The IR images have several features that make it challenging to automatically detect the gland regions:

1. low-contrast between gland and non-gland regions;
2. specular reflections caused from smooth and wet surfaces;

3. inhomogeneous gray-level distributions over regions because of thermal imaging;

4. irregularities in imaged regions of the ocular surface.

Even though there are diversities in imaging conditions, the gland regions have higher reflectance with respect to non-gland regions. Thus, the image regions belonging to a gland have relatively higher brightness with respect to neighbouring non-gland regions. However, because of the above mentioned imaging conditions, conventional methods such as local thresholding are not suitable for partitioning image into gland and non-gland regions.

10 Summary of the invention

The present invention aims to provide automatic processing of images of an ocular region including a plurality of meibomian glands, such as to identify the locations of meibomian glands and/or to obtain numerical data characterising the glands. The numerical data may be used for grading the glands.

15 A first aspect of the invention proposes using an ocular image of a region including meibomian glands to derive a grade indicative of the health of the meibomian glands, by using in the ocular image to obtain one or more numerical parameters characterizing meibomian glands shown in the ocular image; and automatically determining the grade using the one or more numerical parameters.

20 The grade may be used for a screening of patients, such as to identify patients who require more detailed examination. It can also be used to propose treatments to be performed on the patients.

The numerical parameters preferably include at least one of:

(i) at least one parameter characterizing the diversity between the scale parameters of significant features of the image obtained by a scale-space transform; and/or

(ii) at least one parameter obtained by measurement of lines identified in the ocular image and representing respective glands. The measurement may be done of lines individually (e.g. the length of the lines) or relate to the pairs of neighbouring lines (e.g. distances between neighbouring lines).

A second aspect of the invention proposes in general terms that meibomian glands are identified on ocular images using Gabor filtering as a local filtering technique. The parametrization in shape, local spatial support, and orientation of Gabor filtering is particularly suitable for detecting meibomian glands.

5

Brief description of the figures

Embodiments of the invention will now be described, for the sake of example only, with reference to the following figures in which:

Fig. 1, which is composed of Figs. 1(a) to 1(c), shows three captured IR
10 images of ocular surfaces;

Fig. 2, which is composed of Figs. 2(a) to 2(f), shows representations of six Gabor functions;

Fig. 3, is composed of Figs. 3(a) to 3(f), and includes Fig. 3(a) which is portion of Fig. 1(a), and Figs. 3(b)-(f) which illustrate stages of processing the image of Fig. 3(a) using an embodiment of the invention;
15

Fig. 4, is composed of Figs. 4(a) to 4(f), and includes Fig. 4(a) which is portion of Fig. 1(a), and Figs. 4(b)-(f) which illustrate stages of processing the image of Fig. 4(a) using an embodiment of the invention;

Fig. 5, is composed of Figs. 5(a) to 5(f), and includes Fig. 5(a) which is portion of Fig. 1(a), and Figs. 5(b)-(f) which illustrate stages of processing the image of Fig. 5(a) using an embodiment of the invention;
20

Fig. 6, is composed of Figs. 6(a) to 6(f), and includes Fig. 6(a) which is portion of Fig. 1(a), and Figs. 6(b)-(f) which illustrate stages of processing the image of Fig. 6(a) using an embodiment of the invention;

Fig. 7, is composed of Fig. 7(a) and 7(b) which respectively show an ocular image before and after histogram equalisation in a further embodiment of the invention;
25

Fig. 8 is composed of Fig. 8(a) which illustrates Scale Invariant Feature Transform (scale-space) points in a healthy image, and Fig. 8(b) which illustrates scale-space points in an unhealthy image;

Fig. 9 illustrates the distribution of scale invariance and Shannon entropy for a
5 population of healthy and unhealthy images.

Fig. 10, illustrates the extraction of line features in the further embodiment of the invention;

Fig. 11 illustrates a process used by the further embodiment of the invention to extract contiguous lines from pixel clusters;

10 Fig. 12 shows the distribution of the total length of the lines extracted by the process of Fig. 11, and the number of lines, for healthy and unhealthy images;

Fig. 13 is a scatterplot of the lengths of lines, and the standard deviations in the lengths, for healthy and unhealthy images;

15 Fig. 14, which is composed of Figs. 14(a) and 14(b), shows lines superimposed on ocular images, and including both lines representative of glands and spurious lines;

Fig. 15 shows the distribution of the maximum and minimum distances of the ends of the lines from the edge of the images;

Fig. 16 is a variant of Fig. 21 including also points for intermediate images;

20 Fig. 17 is a flow diagram of the first embodiment of the invention; and

Fig. 18 is a flow diagram of a further embodiment of the invention.

Detailed description of the embodiments.

1st embodiment

1.1 Gabor functions

25 The first embodiment of the invention is a method of detecting meibomian glands, making use of the family of two-dimensional (2D) Gabor functions. It is known to use

a Gabor function as a receptive field function of a cell, to model the spatial summation properties of simple cells [1]. A modified parametrization of Gabor functions is used to take into account restrictions found in the experimental data [2, 3]. Suppose there is a light impulse at a point (x, y) on a 2-dimensional visual field Ω (that is $(x, y) \in \Omega \subset \mathbb{R}^2$). The Gabor function is denoted by $G_{\lambda, \theta, \psi}(x, y)$ which is a real valued number (i.e. $G_{\lambda, \theta, \psi}(x, y) \in \mathbb{R}$). The Gabor function is given by [2]:

$$G_{\lambda, \theta, \psi}(x, y) = \exp\left(-\frac{\tilde{x}^2 + \gamma^2 \tilde{y}^2}{2\sigma^2}\right) \cos\left(2\pi \frac{\tilde{x}}{\lambda} + \psi\right) - G_{\lambda, \theta, \psi}^{DC} \quad (1)$$

where

$$\begin{aligned} \tilde{x} &= (x - x_0) \cos(\theta - \pi/2) + (y - y_0) \sin(\theta - \pi/2), \\ \tilde{y} &= -(x - x_0) \sin(\theta - \pi/2) + (y - y_0) \cos(\theta - \pi/2), \end{aligned}$$

and $G_{\lambda, \theta, \psi}^{DC}$ is DC term due to cosine function. The DC term

$$G_{\lambda, \theta, \psi}^{DC} = \iint_{\Omega} G_{\lambda, \theta, \psi}(x, y) dx dy \quad (2)$$

is subtracted from $G_{\lambda, \theta, \psi}$ to remove the bias.

Without loss of generality, it is assumed in the embodiment that Gabor function is centred at the coordinate plane of the receptive field. Thus, x_0 and y_0 are not used to index a receptive field function. The parameters $\sigma, \gamma, \lambda, \theta$ and ψ are explained below.

The size of the receptive field is determined by the standard deviation σ of the Gaussian factor. The parameter γ is in the range 0.23 to 0.92 (i.e. $\gamma \in (0.23, 0.92)$) [2] and is called the spatial aspect ratio. It determines the ellipticity of the receptive field. The value $\gamma = 0.5$ is used in the experimental results below, and, since this value is constant, the parameter γ is not used to index a receptive field function. The parameter λ is the wavelength and $1/\lambda$ is the spatial frequency of the cosine factor. The ratio $\sigma\lambda$ determines the spatial frequency bandwidth, and, therefore, the number of parallel excitatory and inhibitory stripe zones which can be observed in the

receptive field as shown in Fig. 2 (as explained below). The half-response spatial frequency bandwidth $b \in [0.5, 2.5]$ (in octaves) [2] of a linear filter with an impulse response according to Eqn. (1) is the following function of the ratio σ/λ [2]:

$$b = \log_2 \left(\frac{\frac{\sigma}{\lambda} + \frac{1}{\pi} \sqrt{\frac{\ln(2)}{2}}}{\frac{\sigma}{\lambda} - \frac{1}{\pi} \sqrt{\frac{\ln(2)}{2}}} \right) \quad (3)$$

or inversely

$$\frac{\sigma}{\lambda} = \frac{1}{\pi} \sqrt{\frac{\ln(2)}{2} \frac{2^b + 1}{2^b - 1}} \quad (4)$$

- 5 The value $b = 1.0$ is used in the embodiment and, since this value is constant, the parameter σ , which can be computed according to Eqn. (4) for a given λ , is not used to index a receptive field function. The angle parameter $\theta \in [0, \pi)$ determines the preferred orientation from the x-axis in counterclockwise direction. The parameter $\psi \in (-\pi, \pi]$ is a phase offset that determines the symmetry of $G_{\lambda, \theta, \psi}(x, y)$ with respect to the origin: for $\psi = 0$ and $\psi = \pi$ it is symmetric (or even), and for $\psi = -\pi/2$ and $\psi = \pi/2$ it is antisymmetric (or odd); all other cases are asymmetric mixtures.

Fig. 2 is an intensity map of Gabor functions with parameters $\theta = \pi/4$, $\psi = 0.0$ and (a) $\lambda = 10$ pixels; (b) $\lambda = 20$ pixels; (c) $\lambda = 30$ pixels; (d) $\lambda = 40$ pixels; (e) $\lambda = 50$ pixels and (f) $\lambda = 60$ pixels on 256×256 pixels grid, which model the receptive field profile of a simple cell. Gray levels which are lighter and darker than the background indicate zones in which the function takes positive and negative values, respectively. The bright ellipse $\tilde{x}^2 + (\tilde{y})^2 = 4\sigma^2$ specifies the boundary of the (classical) receptive field outside which the function takes negligibly small values.

- 20 Using the above parametrization for Gabor function, one can compute the response $I_{\lambda, \theta, \psi}$ to an input 2D image I as

$$I_{\lambda,\theta,\psi} = I * G_{\lambda,\theta,\psi} \quad (5)$$

where $*$ denotes 2D convolution. Eqn. (5) can be efficiently computed using the Fourier transform ($\mathcal{F}$), i.e., $I_{\lambda,\theta,\psi} = \mathcal{F}^{-1}(\mathcal{F}(I) \mathcal{F}(G_{\lambda,\theta,\psi}))$ where $\mathcal{F}^{-1}$ is the inverse Fourier transform.

1.2 Extracting Features using Gabor Filtering

- 5 Realizations of Gabor functions shown in Fig. 2 can be used to model the local structure of a gland which is surrounded by non-gland regions. That is, the main lobe in the middle represents the gland, and the side lobes on both sides of main lobe represent non-gland regions. The parameter λ can be used as an estimate for the spatial width, and the parameter θ can be used as an estimate to local orientation of
- 10 sub-gland structure. The value $\psi = 0.0$ is used. Without loss of generality, the parameter ψ is not used as index unless otherwise stated.

- The parameter λ takes discrete integer values from a finite set $\{\lambda_i\}$ and can be estimated according to expected spatial width of the consecutive gland and non-gland regions. Meanwhile, it is expected that sub-gland structure can have any
- 15 orientation in between $[0, \pi)$. However, it impossible to test every possible orientation. Thus, the parameter θ is discretized according to:

$$\theta_i = \frac{(i-1)\pi}{N_\theta}, i = 1, 2, \dots, N_\theta \quad (6)$$

where N_θ is the total number of discrete orientations.

- For a rough estimate of correct λ and θ for each pixel, the Gabor filter response is positive over gland region, but it is negative over non-gland regions. In order to
- 20 demonstrate this, a sub-region of IR image from Fig. 1(a) is used, and shown in Fig. 3(a). Below we will explain the steps of using Fig. 3(a), and values $\lambda = 40$ pixels and $N_\theta = 180$ to derive a binarized filter response shown in Fig. 3(f). Fig. 3(b) is a contrast enhanced IR image in which the boundary pixels of the regions of Fig. 3(f) are overlaid, and four pixel locations are labelled as A, B, C, D. Pixels A and B are

on gland regions, and pixel C is on a non-gland region. Pixel D is on the border of a gland region and a non-gland region.

Fig. 3(c) shows the Gabor filter responses of the four pixels shown in Fig. 3(b). The Gabor filter responses for different pixel locations falling into regions of gland and non-gland areas are reported in Fig. 3(c) plotted against the variable θ . The maxima of absolute valued Gabor filter responses of pixels A and B are realized when the sign of Gabor filter response is +1, meanwhile, the sign of the maximum value of absolute valued Gabor filter response is -1 for pixel C. Thus, for a pixel over gland region which is surrounded by non-gland regions, there is a high difference between magnitudes of positive maximum and negative minimum. Similarly, for a pixel over a non-gland region which is surrounded by gland regions, there is a high difference between magnitudes of negative minimum and positive maximum. However, when the pixel resides on a region which has very close proximity between gland and non-gland regions, e.g., pixel D on Fig. 3(b), the difference between magnitudes of positive maximum and negative minimum is small, as shown in Fig. 3(c). It is also clear that the sign of average Gabor response is +1 over gland regions, and is -1 over non-gland regions.

For a given λ , the mean Gabor response $\hat{I}_\lambda$ is computed as follows:

$$\hat{I}_\lambda = \frac{\int_0^\pi I_{\lambda,\theta} d\theta}{\int_0^\pi d\theta} = \frac{1}{\pi} \int_0^\pi I_{\lambda,\theta} d\theta, \quad (7)$$

which is approximated as

$$\hat{I}_\lambda = \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} I_{\lambda,\theta_i}. \quad (8)$$

In Fig. 3(d) the average Gabor filter response $\hat{I}_{\lambda=40}$ of the input image Fig. 3(a) computed according to Eqn. (8) is shown for $\lambda = 40$ pixels and $N_\theta = 180$. The corresponding surface plot is given in Fig. 3(e). It is clear from Fig. 3(d) and 3(e) that the filter gives high positive responses on gland regions, while it produces low

negative responses on non-gland regions. Thus using this observation, one can easily segment out gland regions from non-gland regions using the sign of filter response, i.e.,

$$\hat{B}_\lambda(x, y) = H\left(\hat{I}_\lambda(x, y)\right), \quad (9)$$

where $H(a) \in \{0, 1\}$ is a binary function defined as

$$H(a) = \begin{cases} 1, & \text{if } a \geq 0; \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

- 5 In Fig. 3(f) the binarized Gabor filter response $\hat{B}_\lambda$ according to Eqn. (9) is shown. Overall, the segmentation result looks satisfactory. However, it can be observed that at pixel D two glands regions are merged together where the separation between two glands is distinct in the surface plot shown in Fig. 3(e). This can be due to two reasons: 1) the value of λ is too large to resolve the signal separation; and/or 2) the pixel D falls into an uncertain region where there is not enough information to
10 separate two regions or the region belongs to some other parts of the ocular surface. In latter case, further analysis is required. However, the former case can be resolved by analysing the image at different values of λ . For example in Fig. 4, the same input image shown in Fig. 3(a) is analysed using $\lambda = 20$ pixels and $N_\theta = 180$. Fig. 4(a) is
15 the same as Fig. 3(a), and each other part of Fig. 4 corresponds to a respective image in Fig. 3. It is clear that the pixel D is segmented correctly at the expense of incorrect segmentations in other regions.

Corresponding results for $\lambda = 30$ pixels and $N_\theta = 180$ are shown in Fig. 5, in which each image corresponds to the respective image in Fig. 3.

- 20 The results obtained by exploiting different values of λ show various trade-offs yielded between spatial-detail preservation and noise reduction. In particular, images with a lower value of λ are more subject to noise interference, while preserving more details of image content. On the contrary, images with a higher value of λ are less susceptible to noise interference, but sacrificing more degradations on image details.

The average Gabor filter responses for each distinct value of $\lambda \in \{\lambda_i\}$ is used in vector representation $\mathbf{f}_{x,y}$ for each pixel (x, y) of input image as

$$\begin{aligned} \mathbf{f}_{x,y} &= [\mathbf{f}_{x,y}(1), \dots, \mathbf{f}_{x,y}(i), \dots, \mathbf{f}_{x,y}(N_\lambda)]; \\ &= \frac{[I_{\lambda_1}(x,y), \dots, I_{\lambda_i}(x,y), \dots, I_{\lambda_{N_\lambda}}(x,y)]}{\sqrt{\sum_{i=1}^{N_\lambda} I_{\lambda_i}^2(x,y)}}, \end{aligned} \quad (11)$$

where N_λ is the cardinality of the set $\{\lambda_i\}$, i is a positive integer (i.e. $i \in \mathbb{Z}^+$), and

$$\mathbf{f}_{x,y}(i) = \frac{I_{\lambda_i}(x,y)}{\sqrt{\sum_{k=1}^{N_\lambda} I_{\lambda_k}^2(x,y)}}. \quad (12)$$

The denominator in Eqn. (12) is used to compensate the fluctuations due to illumination differences over different parts of the image. In Fig. 6, Gabor filter responses according to Eqn. (12) are shown for $N_\lambda = 16$ and $N_\theta = 180$. Fig. 6(a) is the same as Fig. 3(a), 4(a) and 5(a). Below we will explain the steps of using Fig. 6(a) to derive a binarized filter response shown in Fig. 6(f). Fig. 6(b) is a contrast enhanced IR image in which the boundary pixels of the regions of Fig. 6(f) are overlaid, and the four pixel locations are labelled as A, B, C, D are as in Figs. 3 to 5. Fig. 6(c) shows the Gabor filter responses of the four pixels shown in Fig. 6(b) according to Eqn (12). The Gabor filter responses for different pixel locations falling into regions of gland and non-gland areas are plotted in Fig. 6(c). From the feature vectors shown in Fig. 6(c), the feature vectors are positive on gland regions, and negative in non-gland regions. However, they fluctuate between negative and positive values for pixel D. The average feature $\hat{F}$ is computed as

$$\hat{F}(x, y) = \frac{1}{N_\lambda} \sum_{i=1}^{N_\lambda} \mathbf{f}_{x,y}(i). \quad (13)$$

which is depicted in Fig. 6(d) and (e) where the discrimination between gland and non-gland regions are clear. Using Eqn. (13), the final binary image is computed according to

$$\hat{B}(x, y) = H\left(\hat{F}(x, y)\right). \quad (14)$$

where $H(\cdot)$ is defined in Eqn. (10). The result of binarization according to Eqn. (14) is shown in Fig. 6(f) and the boundaries are overlaid onto input image in Fig 6(b). It is clear that merging information from different values of λ improves the discrimination between gland and non-gland regions.

- 5 Thus, the steps of the first embodiment are as shown in Fig. 17. In step 1 the value of λ is initialised (i.e. i is set to first value, so as to set λ_i). In step 2, θ is initialised. In step 3, the values of λ and θ are used to perform a Gabor filter transform. This step is repeated for each of the possible values of θ , and the result is used in step 4 to form the value of I_λ for this value of λ . The result is iterated over the possible values
 10 of λ , and the result is used to form $\hat{B}$ in step 5 using Eqns. (13) and (14).

2nd embodiment

The second embodiment aims to provide a way of grading a subject, i.e. allotting him into one of at least two categories, such as "healthy", "unhealthy" or "intermediate".

- The overall method of the second embodiment is illustrated in Fig. 18. A single
 15 ocular image is used to obtain one or more numerical parameters ("features") indicative of whether the image is healthy or not. Note that not all the numerical parameters described below may be collected in realisations of the embodiment, but preferably more than one parameter is collected, and in this case the numerical parameters are combined by an adaptive learning system (such as a support vector
 20 machine (SVM) which has been subject to supervised learning), to generate an output indicative of whether the image is healthy or not.

2.1 Enhancing Contrast Using Histogram Equalization

- The original images have poor contrast. To improve contrast, we applied a standard technique called Histogram Equalization. The original image is shown in Fig. 7(a)
 25 and the image with improved contrast is shown in Fig 7(b) for comparison. The operations in the second embodiment are performed on the images after Histogram Equalization.

2.2 Scale-space-Shannon Entropy Feature

2.2.1 Scale-space transform

The second embodiment employs a feature called the Scale-Space-Shannon Entropy feature to distinguish a healthy from an unhealthy image. This concept is adapted from a well-known method called Scale Invariant Feature Transform (SIFT) described in [4]. In short, in step 20 of Fig. 19, the embodiment locates keypoints on an image, called scale-space points. Each scale-space point is represented by a vector with 3 elements (x, y, s) (note that by contrast the SIFT transform uses a further 129 elements which are not employed in the embodiment). x and y are the Cartesian coordinates of the scale-space point on the image, s is called its *scale*. The scale-space points are found by the following "scale-space transform":

- convolving the (x, y) ocular image with Gaussian filters at different distance scales s ;
- seeking maxima (x, y, s) of the difference between subsequent pairs of the images with different s to form candidate keypoints; and
- rejecting candidate keypoints for which the contrast with next-neighbour points is below a threshold.

Fig. 8 shows how the scale-space points look on two images: a healthy image (Fig. 8(a)) and an unhealthy image (Fig. 8(b)). Each scale-space point is represented as a circle, and the horizontal bar of the circle (i.e. the radius) indicates its scale s .

The embodiment employs the observation that, as shown in Fig. 8, within a local region (shown by the boxes), the circles of healthy images are of similar sizes, whereas the circles of unhealthy images are of very different sizes. This is because in healthy images, there are evenly-spaced strips of similar thickness (i.e. the glands), and the scale-space transform picks up this pattern. Unhealthy images, on the other hand, do not have this pattern (i.e. no glands), and the sizes of the circles are therefore of very different magnitudes.

In step 21 of Fig. 18, the embodiment generates a numerical measure of the disparity of the sizes of the circles, and the embodiment makes use of the fact that the local distribution of scales is uniform for a healthy image and non-uniform for an unhealthy one to distinguish between the two classes. One mathematical function
 5 which can measure this uniformity is the well-known Shannon Entropy, which we will now discuss.

2.2.2 Shannon Entropy

Shannon entropy is defined as

$$S = - \sum_{i=1}^n p_i \ln p_i, \quad (15)$$

where p_i is the probability of event i . The p_i 's must be normalized, i.e. $\sum_{i=1}^n p_i = 1$.

10 One important property of the Shannon Entropy that we use is that it is maximized if and only if p_i is a uniform distribution, i.e. $p_i = 1/n$. For the proof, refer to [5].

The scale of the SIFT points can be related to the probability distribution in the following way. First, choose a scale-space point and consider its n nearest neighboring scale-space points. We define the 'probability' of the i -th nearest
 15 neighbor with respect to this 'central' scale-space point as

$$p_i = \frac{(s_i)^2}{\sum_{i=1}^n (s_i)^2} \quad (16)$$

where $i=1, \dots, n$ labels the n nearest-neighbouring scale-space points and s_i is the scale of the i -th scale-space point.

The denominator in Eqn. (16) ensures that $0 < p_i < 1$, and that the distribution is normalized. The meaning of p_i is the ratio of the area of the circle of the i th neighbor
 20 to the total area of all the neighbors.

Referring back to Fig. 8, we see that for the box in Fig. 8(a), all the circles are roughly the same size, so p_i will also be roughly the same, and hence the entropy computed

within that blue box should be close to maximal. For Fig. 8(b), on the other hand, because the circles are very non-uniform, p_i computed according to Eq. (16) is also very non-uniform, and so the entropy computed within that blue box will be low compared to (a).

5 2.2.3 Algorithm to Compute Shannon Entropy of an Image

To compare the entropies of two different images, we need to compute the Shannon Entropy for an entire image. The algorithm to do this is as follows.

1. Obtain all the scale-space points for an image (using standard techniques). Say the total number of scale-space points is M . Denote a scale-space point as α , where $\alpha = 1, \dots, M$.
2. For an α , identify its n nearest neighbors (usually $n = 20$). (In cases in which the n nearest neighbours are not unambiguously defined (e.g. because there are 3 scale-space points exactly the same distance away and 19 scale-space points closer; or to put this more generally, if for the smallest distance d such that there are at least n points no further than d from α , the number of scale-space points no further than d from α is m which is greater than n), the algorithm can randomly take n of these m points, or alternatively use all m of these points)
3. If $\{s_1^\alpha, \dots, s_n^\alpha\}$ are the scales of these n nearest neighbours, let $s^\alpha = \sum_{i=1}^n (s_i^\alpha)^2$, and let $p_i^\alpha = (s_i^\alpha)^2 / s^\alpha$.
4. Compute the entropy for α : $S^\alpha = -\sum_{i=1}^n p_i^\alpha \ln p_i^\alpha$
5. Repeat Steps 1 to 3 for all scale-space points, i.e. $\alpha = 1, \dots, M$.
6. The Shannon Entropy for the entire image is the average

$$\bar{S} = \frac{1}{M} \sum_{\alpha=1}^M S^\alpha \quad (17)$$

2.2.4 Results

- 25 To see whether the scale-space-Shannon Entropy feature can distinguish between healthy and unhealthy meibography images, we manually graded some healthy and unhealthy images, and computed their entropy S . The results are shown in Fig. 9.

Each image is represented as a point. The horizontal axis represents entropy S . The vertical axis is included only for ease of visualisation. Healthy images are plotted as lighter dots, and unhealthy images as respective darker dots. The result shows that the healthy cluster is quite well separated from the unhealthy cluster along the Shannon Entropy dimension.

2.3 Line Features

The embodiment uses a further method to extract other features, which we call Line Features, from the images. As explained, the most salient characteristic of a healthy image are the vertical gland patterns, shown in the top left panel of Figure 10. The embodiment obtains clusters of pixels indicative of these glands, as shown in the two bottom panels of Figure 10. The procedure can be summarized as follows.

1. Extracting pixels that lie along the bright and dark line regions (i.e. gland patterns). This is step 22 of Fig. 18.
2. Grouping the pixels into 'primitive' clusters that resemble line patterns. This is step 23 of Fig. 18.
3. Morphological operations on each cluster to form a line from the cluster. This is step 24 of Fig. 18.
4. Obtaining numerical properties of the lines (e.g. their length and/or curvature) as features for classification. This is step 25 of Fig. 18.

In the following, we shall describe each of the above steps.

2.3.1 Extracting Local Minima and Maxima Pixels

To extract the local minima, we first smooth the entire image (implemented using the cvSmooth algorithm from the well-known OpenCV library of programming functions) so that the values of the intensity change smoothly from one pixel to another. Then, using the intensity as a kind of 'potential energy surface', we can seek minima of the surface.

One way to do this is to perform gradient descent to reach the minima of the surface. Specifically, the procedure is as follows.

1. Consider a horizontal cross section of the image, i.e. one row. Denote the intensity at the i th pixel as $I[i]$.
2. Consider a 'particle' located at the i th pixel.
3. Compute the 'force' on it as

$$f[i] = -(I[i+1] - I[i-1]). \quad (18)$$

- 5 (which is just a simplified version of the gradient formula $-\partial V/\partial x$.) If $f[i] > 0$, move the particle to $i+1$; if $f[i] < 0$, move it to $i-1$; if $f[i] = 0$, then it stays at i . Note that since intensities are generally quoted in integer values (e.g. integers in the range 0 to 255) one would expect $f[i] = 0$ at minima.
4. Repeat Step 3, until the particle stops moving, i.e. $f = 0$. Denote the location of the particle as $\hat{i}$.
 - 10 5. Repeat Steps 2 and 3, starting the particle from every pixel along the row, obtaining an $\hat{i}$ for each starting pixel i . The set of all $\hat{i}$'s are the minima for the row.
 6. Proceed to the next row and repeat Steps 2 to 5.
- 15 To obtain the maxima, one simply repeats the above but using $-I$ instead of I for the pixel intensity.

An alternative algorithm is to scan along the row, and make a list of those points/pixels i that satisfy the following conditions: for maxima: $I[i-1] < I[i] > I[i+1]$; for minima: $I[i-1] > I[i] < I[i+1]$.

20 2.3.2 Clustering the Pixels

- After obtaining all the minima and maxima points, we need to cluster them so that, to a first approximation, each cluster has a line-like shape. This will facilitate the next stage of converting the cluster into a contiguous line. To do our clustering, we use a well-known algorithm called First In First Out (FIFO). For specificity, we use the
- 25 threshold of 10 pixels, which means that only pixels within less than 10 pixels of some other pixels in a cluster are grouped together. After this procedure, the minima (maxima) are grouped as shown in the bottom two panels of Fig. 10. Each of these

cluster are then transformed into a contiguous curve using the method described in the next section.

2.3.3 Transforming a Cluster into a Contiguous Curve

At this stage, each cluster resembles a line, but it is still not useful because it may be broken, and it is not one pixel thick. Here, we present an algorithm, based on well-known methods, to transform a cluster into a curve. While each of the methods used by the embodiment is known as unrelated method, the embodiment combine these methods to achieve the purpose of converting a cluster of pixels into a contiguous curve.

The sub-steps of step 24 in Fig. 18 are illustrated in Fig. 11, and are described below.

1. Choose one cluster and put it into an image. Specifically, we set all the pixels belonging to that cluster as foreground (black in Fig. 11), and all other pixels as background (i.e. white). This is shown in the first panel of Fig. 11.
2. Next, we apply cvDilate (which is another algorithm available in OpenCV) to thicken the cluster by one pixel. The purpose is to merge all the pixels into one connected piece. After one application of cvDilate, we check to see if the cluster now consists of a single connected component. If yes, we proceed to the next step, otherwise, we apply cvDilate again until one single connected component is obtained.
3. The previous step may produce a connected component containing 'islands' of background, highlighted by the circle in the second panel of Fig. 11. These islands must be eliminated because they will give rise to 'loops' in the contiguous line in the later steps. To eliminate all them, we use cvFloodFill (which is another algorithm available in OpenCV) to fill out the background first, revealing the locations of these islands as the remaining white pixels. We then go back to the pre-cvFloodFill image and set these islands to foreground (i.e. black). This produces a connected component which does not contain these problematic islands.
4. We then use a standard thinning algorithm [6] to thin the component until one pixel thick.

5. The connected component is now a 'tree' with many branches. The embodiment prunes away all the side branches. To do this, we first locate and count the number of terminal points (i.e. end points) on the tree. We then use a standard pruning algorithm [7] to prune away all side branches with pruning factor 1, i.e. branches which are 1 pixel long are eliminated. We then count the number of end points again, and if it is not 2, we increase the pruning factor by 1 and prune again. The procedure is repeated until only two end points are left, meaning that the tree has no more branches left.
6. This completes the procedure for transforming a cluster into a contiguous line. We then go repeat Steps 1 to 5 for another cluster.

2.3.4 Selecting Features, and Classification of Healthy and Unhealthy Images

We now present features (i.e. numerical parameters) derived in step 25 of Fig. 18 from the lines obtained to help us distinguish between the healthy and unhealthy images. The features are:

1. Number of lines.
2. Total length of lines.

The following all alternative features which may be used.

3. Potential energy.
4. Left-Right distance.
5. Twistedness.

No. of Lines vs Total Length As the first two features, we choose the 'number of lines' and 'total length of the lines'. The idea is explained graphically in Fig. 12. For a healthy image, there are fewer lines, and all the lines are long; for an unhealthy image, there are more lines, and the lines are short. If we plot these two features, the healthy and unhealthy images can be separated, as shown in the bottom panels of Fig. 12.

Potential Energy The next feature is called 'potential energy'. We first find the nearest neighbor of a minimum (maximum) point. We then compute its potential energy. The potential energy is zero everywhere except at the distance of around 50

pixels ($\approx$ spacing between strips) where it becomes negative. Hence, for healthy images, we expect the total potential energy of an image to be negative, whereas for unhealthy images we expect it to be close to zero.

Left-Right Distance The next feature is called Left-Right Distance. We start off from a point, and move in the two directions perpendicular to the tangent at that point. For each direction, we compute the distance where the embodiment first encounters another point (d_1 and d_2 in the figure). The distribution of points on the d_1 - d_2 plane can be used to construct a histogram. The first two components of the histogram are good features to separate healthy and unhealthy images.

Twistedness Our last feature is called Twistedness, and it is based on the observation that the lines for healthy images are less twisted than those for unhealthy images. To quantify twistedness, we define it as the ratio of the distance between the ends of the line ("straight length") and the length of the line ("arc length"). For each image, a histogram can be computed for the distribution of twistedness of its lines. We then further coarse grain the histogram into two bins, and use these as the feature space.

2.4 Variations in the second embodiment

The features discussed above are not the only features which can be derived in step 25 from the lines obtained in step 24. Two further possible features are the average of the length of all the lines in an image, and the standard deviation of length of the lines. Fig. 13 shows the scatter plot of the images based on these two features. Each point represents an image. The circle-shaped and diamond-shaped points represent the healthy and unhealthy images respectively. It is clear that using standard techniques like Support Vector Machines, the two classes of images can be classified. This method has been reported in our recent publication in Journal of Biomedical Optics 17(8) 0860008, (2012).

Furthermore, many variations are possible in the technique used to obtain the lines, and improvements can be made. Although the algorithm presented above for extracting line features from the image is effective for classification of the extreme

cases of healthy and unhealthy images, it may be less effective to assess the intermediate stages, and it is valuable to minimise noise. Noise here means spurious lines that are extracted by the algorithm of Fig. 18 but are not due to meibomian glands. The majority of these lines are due to eye lashes or dark shadows at the periphery of the images. These are removed in order to remove noise from the subsequent statistics that we will compute from the lines. An example is shown in Fig. 14(a). Lines extracted by the method of Fig. 18 are marked. Lines marked X are due to the meibomian glands, but the lines marked Y near the edge are spurious and due to inhomogeneity in the pixel intensity, and are preferably excluded before step 25 is performed, i.e. excluded from subsequent calculation of statistics based on the lines. In the following, we describe an algorithm for automatically detecting the spurious lines Y.

The spurious lines Y have two important properties. They are much shorter than the majority of the lines, and they are close to the edges. However, it would not be appropriate to exclude all lines which are short because they are characteristic of unhealthy and intermediate images, as shown in Fig. 14(b). We see that due to the breaking up of the gland patterns in the center, there are many short broken lines in the center of the image. These must be retained.

There are two stages to remove the spurious lines. In the first stage, we process each individual line in the following way.

- (1) Label the top, bottom, left, and right edges of the images (as, for instance, 1,2,3,4 respectively). For every pixel of a line, check which side the pixel is closest to, and determine the distance of the pixel to this closest side. Call this number d .
- (2) For each line, get the pixel with the smallest and largest d . Call them $d_{\min}$ and $d_{\max}$ respectively.
- (3) Compute the $d_{\min}$ and $d_{\max}$ for all the lines on the images. Plot each line as a point in a two dimensional space where the x-axis is $d_{\max}-d_{\min}$, and the y-axis is $d_{\max}$. The resulting scatter plot is shown in Fig. 15.

All the lines fall into one of three sectors in the 2-D plot. The lines which correspond to gland lines will lie in the top-right hand region. This is because they usually contain a pixel which is close to an edge, and hence a small $d_{\min}$, but also contain

a pixel which is near the center, and hence a large $d_{\max}$. As a result, lines associated with glands will have large $d_{\max}-d_{\min}$ and $d_{\max}$.

A second category of lines belong to those which fall in the top left hand region.

These are remnants of broken-up gland lines. They are short and lie near the center.

- 5 As they lie near the center of an image, their $d_{\max}$ is large. But because they are short and near the center, the pixel with $d_{\min}$ is usually also in the proximity to the pixel with $d_{\max}$, and $d_{\min}$ is approximately equal to $d_{\max}$, giving a small $d_{\max}-d_{\min}$. This explains why they lie in the top left hand corner. These lines should also be included when step 25 is performed.

- 10 The spurious lines Y are those that lie close to the origin of the scatter plot. They are short and hence every pixel along the line will be close to the edge, hence $d_{\max}$ will be small, and so will $d_{\max}-d_{\min}$.

To select the spurious lines, we use a thresholding procedure. Any line that has $d_{\max}-d_{\min}<50$ and $d_{\max}<50$ will be classified as spurious and hence removed

- 15 before step 25 is performed.

2.5 Industrial applicability

- Fig. 13 showed that healthy and unhealthy images can be well-distinguished using the average of the length of all the lines in an image and the standard deviation of the lines. This, and the other experimental results given above, imply that the second
- 20 embodiment can used as a successful tool for identifying subjects in these categories, even using a single ocular images. The subjects found to have unhealthy eyes can be subjected to further examination, or treated. It is known to treat patients with meibomian gland dysfunction with a warm compress with a hot towel, an
- 25 eyemask, or a specialised heating device called blephasteam. In addition, anti-inflammatory medications such as doxycyclines, azithromycin and cyclosporin may be helpful. In addition, patients may be started on antibiotic steroid ointments topically. Off the counter lubricants, especially those containing lipids may be given to replenish the tear lipids, since there may be abnormal tear lipid layer in patients
- 30 with this condition. Recent advances in this condition include mechanical probing of meibomian glands with or without intraglandular injection of steroids, as well as the

in-office thermopulsation treatment called Lipiflow. As many of these treatments may be prone to poor compliance, or involve complications, or are expensive, it is helpful to have an objective means of diagnosing and evaluating meibomian gland dysfunction.

- 5 Fig. 16, which uses the same two features as Fig. 13 but also includes points for intermediate images shows that for these two features the intermediate grades overlap significantly with the healthy and unhealthy grades. Thus, it is desirable to use other features, such as the ones described above, to help distinguish the intermediate images. For example, it is intended to use the classifier described
- 10 above for separating the healthy and unhealthy images to first give a preliminary classification of images into two classes (i) healthy/intermediate images or (ii) unhealthy/intermediate images. Specific features that distinguish between healthy/intermediate images can then be used to further separate the healthy from the intermediate class in the healthy/intermediate category. Similarly for the
- 15 unhealthy/intermediate category.

References

- [1] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1160–1169, Jul 1985.
- 20 [2] N. Petkov and P. Kruizinga, "Computational models of visual neurons specialised in the detection of periodic and aperiodic oriented visual stimuli: bar and grating cells," *Biological Cybernetics*, vol. 76, pp. 83–96, 1997.
- [3] P. Kruizinga and N. Petkov, "Nonlinear operator for oriented texture," *IEEE Transactions on Image Processing*, vol. 8, no. 10, pp. 1395–1407, Oct 1999.
- 25 [4] D. G. Lowe, "Object Recognition from Local Scale-Invariant Features," *The Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, pp. 1150–1157, 1999.
- [5] A. I. Khinchin, *Mathematical Foundations of Information Theory*. Dover Publications, 1957.
- 30

- [6] L. Lam, S. Lee, and C. Suen, "Thinning methodologies-a comprehensive survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 869–885, 1992.
- 5 [7] A. Niemisto, V. Dunmire, O. Yli-Harja, W. Zhang, and I. Shmulevich, "Robust quantification of in vitro angiogenesis through image analysis." *IEEE Transactions on Medical Imaging*, vol. 24, no. 4, pp. 549–553, 2005.
- 10 [8] K. K. Nichols, G. N. Foulks, A. J. Bron, B. J. Glasgow, M. Dogru, K. Tsubota, M. A. Lemp, and D. A. Sullivan, The International Workshop on Meibomian Gland Dysfunction: Executive Summary, *Ophthalmol. Vis. Sci.* March 30, 2011 vol. 52 no. 4 1922-1929.

Claims

1. A method performed by a computer apparatus, for using an ocular image of a region including meibomian glands to derive a grade indicative of the health of the meibomian glands, the method including:
 - 5 (i) automatically obtaining one or more numerical parameters characterizing meibomian glands shown in the ocular image;
 - (ii) automatically determining the grade using the one of more numerical parameters.
- 10 2. A method according to claim 1 in which said step of obtaining the one or more numerical parameters includes:

generating keypoints in the image, each keypoint being associated with a respective distance scale value;

obtaining one of said numerical parameters as a parameter indicative of the

- 15 disparity of the scale values of the keypoints.
- 3. A method according to claim 2 in which said numerical parameter is calculated as an average over the keypoints of respective numerical value S , the numerical value S for each keypoint being calculated based on a sub-set of the other
- 20 keypoints which are proximate the keypoint.
- 4. A method according to claim 3 in which the numerical value S is calculated according to the expression

$$S = - \sum_{i=1}^n p_i \ln p_i$$

where $i=1, \dots, n$ labels the n keypoints of the subset, and p_i is given by

$$p_i = \frac{(s_i)^2}{\sum_{i=1}^n (s_i)^2}$$

where the s_i are the scale values of the n keypoints.

5. A method according to any preceding claim in which at least one of said numerical features is derived by:

5 obtaining lines in the ocular images indicative of the respective meibomian glands;

obtaining the at least one of the numerical features by measurement made using the lines.

10 6. A method according to claim 5 in which in which the step of obtaining the lines in the images comprises:

extracting pixels that lie along bright and dark line regions of the image; and

grouping the pixels into clusters.

15 7. A method according to claim 6 in which the step of obtaining the lines further comprises morphological operations on each cluster to form a line from the cluster.

8. A method according to any of claims 5 to 7 including a line exclusion step of identifying ones of said lines which are not associated with glands, and removing them from consideration.

20 9. A method according to claim 8 in said line exclusion step is performed by determining for each line the minimum distance $d_{\min}$ to a side of the image, and the maximum distance $d_{\max}$ to a side of the image, and identifying ones of said lines which are not associated with glands using said distances.

25 10. A method according to claim 9 in which the lines which are not associated with glands are identified as those for which $d_{\max}$ and $d_{\max} - d_{\min}$ are below respective thresholds.

11. A method according to any of claims 5 to 10 in which the numerical parameters obtained using the lines, comprise any one or more of:

- (i) the number of lines;
 - (ii) the total length of lines;
 - (iii) the average length of the lines;
 - (iv) the standard deviation of the lengths of the lines;
 - 5 (v) a value obtained by, for a plurality of the lines, finding the sum along each line of a value obtained as a function of the distance of points on the line to points on another line;
 - (vi) a value obtained by, for a plurality of the lines, finding the sum along each line of a value obtained by measuring at least one distance, perpendicular to the
 - 10 tangent of the line, from the line to another said line; or
 - (vii) a value indicative of the ratio of the distance between the ends of the lines and the length of the corresponding lines.
12. A method of segmenting an ocular image of a region including meibomian glands, the method including:
- 15 at each of a plurality of locations in the image:
 - (i) subjecting the image to a Gabor function transform centred at the location, and characterized by at least a scale factor λ and a direction θ within the image;
 - (ii) summing the Gabor function transform over values of θ , to form an intensity value $\hat{I}_\lambda$; and
 - 20 (iii) using $\hat{I}_\lambda$ to perform a thresholding step, to form a binary value representative of whether the corresponding location corresponds to the position of a meibomian gland.
13. A method according to claim 12 in which in step (iii) the value $\hat{I}_\lambda$ is summed over a plurality of values of λ , and the result is thresholded.
- 25 14. A computer apparatus for analysing an ocular image, the apparatus including a processor, and a data storage device storing program instructions operative when

performed by the processor to cause the processor to perform a method according to any of claims 1 to 13.

15. A computer program product comprising program instructions operative when performed by a processor to cause the processor to perform a method according to
5 any of claims 1 to 13.