

Automatic Detection of Pre-Ocular Tear Film Break-Up Sequence in Dry Eyes

Tamir Yedidya and Richard Hartley

The Australian National University and National ICT Australia *

{tamir.yedidya, richard.hartley}@rsise.anu.edu.au

Jean-Pierre Guillon

Faculty of Medicine and Health Sciences, Lions Eye Institute, Australia

jean.pierre@lei.org.au

Abstract

Dry Eye Syndrome is a common disease in the western world, with effects from uncomfortable itchiness to permanent damage to the ocular surface. Nevertheless, there is still no objective test that provides reliable results. We extend our method [13] for the automated detection of dryness signs to include the break up time (BUT), analysis of the degree of thinning of the tear film and detection of meniscus induced dryness, the last two have not been previously addressed. Our motivation is to help the clinician to automatically detect and analyze various signs related to dry eyes. The method has been tested on over 100 videos taken from 30 different patients. When compared to an analysis done by a specialized optometrist, our method is demonstrated to provide an accurate estimation for the BUT and the extent of the disease.

1 Introduction

One of the roles of the tear film is the maintenance of corneal epithelial integrity and transparency which is achieved by keeping the ocular surface continuously moist. The pre-ocular tear film in humans does not remain stable for long periods of time [6]. When blinking is prevented, the tear film ruptures and dry spots appear over the cornea [7].

The Fluorescein Break Up Time (FBUT) test was designed by Norn [9]. A moistened fluorescein strip is applied to the conjunctiva and after a few blinks to

spread the fluorescein evenly, the tear film is viewed with the help of a yellow filter in front of a slit-lamp biomicroscope. When a dark area appears, the lack of fluorescence represents the rupture of the tear film and the time elapsed since the last blink is recorded as break-up time (BUT). The shorter the BUT, the more likely the diagnosis of dry eye [12]. The degree of blackness is related to the depth of the breakup. The deeper the break, the greater the chances of ocular surface damage. If the eyes are kept open, the area of the break will increase in size and breaks may appear in new areas over the cornea. This is the test of tear film stability most commonly used by clinicians as it is minimally invasive [2].

After a break has happened, it is important to further analyze the break area. The location and depth of the original and successive breaks give clinical indications of the cause of the break and are helpful in determining what treatment to choose [1]. This is not possible with the current clinical routines due to the subjectiveness of the grading methods. With the development of the Eye-Scan system [13], an easy to operate, multi purpose system of video recording of anterior ocular structures is available.

The exposed tear volume is divided into the pre-ocular part, a thin layer covering the cornea and a marginal part, called lid tear meniscus or tear reservoir and is situated along the upper and lower lids. The most critical part of the tear film is the junction between the meniscus and the pre-ocular tear. At that point, the surface tension forces make the formation of a continuous film nearly impossible [5]. The result is a line of minimal thickness often referred to as the black line because of its appearance when fluorescein is instilled in the eye. That junction of minimal thickness corresponds to the zone of tear film greatest instability.

*National ICT Australia is funded by the Australian Government's Backing Australia's Ability initiative, in part through the Australia Research Council




Figure 1. Samples of eye images after instilling fluorescein. (a) Immediately after a blink (b) just before the next blink. The darker areas are the dry eye areas which form through the sequence. They are pointed out by arrows. The darker the area is, the drier it became. One sees that dryness forms at different locations with varying severity in the iris area. The upper and lower bright curves are the tear menisci.

Automatic methods for finding the dry eye areas are sparse in the current literature. However, there are some existing approaches for locating the iris as part of other application mainly for iris recognition. Usually, assuming its shape is a perfect circle, the methods mostly use circle fitting algorithms to first locate the pupil and then the iris. Daugman [3] focused on iris recognition, and finds the pupil-iris and iris-sclera borders by searching over all circles in different radii for the ones that gives the maximum contour integral derivative. Ma et al [8] also focus on iris recognition and first locate the darker pupil and then the iris using Canny edge detector and Hough transform. However, Hough transform tends to be very slow and needs a few iterations to find the correct radius. Ritter and el [10] detect the pupil-iris border in slit lamp images in order to register the images. Their method is based on active contours that iteratively grow to fit to the iris border. The acting forces use differences in intensity between the iris and the pupil and aim for a perfect polygon. However, as seen in Fig. 1, the pupil is not visible in our videos and the boundaries between the iris and the sclera are fuzzy due to the fluorescein spreading. Therefore, different solutions are needed.

In this paper, we extend our previous work [13] in several ways:

- Detection of the break up time (BUT).
- Separate detection of meniscus related dryness (black line).
- Improved detection of the eyelids.

- Improved detection of the iris, when strong edges and noise are found on the conjunctiva (see Fig. 1(b)).
- Adapting the function which calculates the degree of thinning to handle large differences in illumination between videos and parts of the same image.
- Analyzing dryness to include the area and location of the break.

2 Algorithm

The dry areas always appear as darker areas in the fluorescent image, however because of changing light conditions and the amount of fluorescein that has been instilled, it is insufficient simply to use some kind of threshold on the last frame. Moreover, in order to find the BUT, the whole video has to be scanned to pin-point the exact time when a certain area became dry. Also the camera is hand held resulting in movement and change of scale of the iris. To overcome these problems, our algorithm is based on three main steps: 1) Detection of the iris and eyelids in the first frame 2) Alignment of the iris in the rest of the sequence and 3) Scanning the aligned video to find the dry areas and BUT.

2.1 Detection of the iris and eyelids

The images of interest in the video¹ are those between consecutive blinks. To that end, we first find all the blinks and half-blinks and treat each sequence individually. Blinks are detected when consecutive images have big differences in intensity. Given an image $I(x, y)$ after a blink, we create an edge map $E(x, y)$ of the image using the Canny edge detector. We found that Canny produces superior results to the Sobel edge detector that was used before, mainly due to the smoothing and the removal of noisy disconnected areas, see Figs. 2(a&b). Then we create 3 threshold images: of the iris, $I_{\text{iris}}(x, y) = (E(x, y) > T_1)$, and the lower and upper eyelids, $I_{\text{low}}(x, y) = (E(x, y) > T_2)$ and $I_{\text{up}}(x, y) = (E(x, y) > T_3)$. We use $T_1 \leq T_2 \leq T_3$ as the edges of the eyelids are usually stronger than the iris's borders, which are fuzzy due to the fluorescein spreading. The parameters for the Canny have been learnt using a subset of the total videos. All videos were taken using the Eye-Scan camera, so the amount of light emitted and the angle of view were

¹The captured videos are of resolution 352x288. The expected range of radii for the iris is between 90 to 130 pixels for most patients.




Figure 2. Results for fitting the iris and eyelids: Threshold image using (a) Sobel (b) Canny. The noise is much less noticeable, but the iris circle has some missing gaps. Also notice the double edges for the eyelids. (c) Pixels chosen by LM for the iris fitting (in white) (d) The curves fitted for the eyelids and the iris.

controlled. However, the strength of the edges still varies among patients and artifacts on the conjunctiva can create strong edges. Therefore, we start with the learnt thresholds and adaptively update them, if the detection of the iris or eyelids fails.

Although our aim is to perform a fitting to the iris and use it later for the alignment, it is impractical to do so directly due to the high ratio between the inliers (the iris pixels) and the outliers (everything else), see Fig. 2(b). In this work we fit a polynomial of degree 2 to the eyelids. In most cases it is an adequate fitting, as we are mainly interested in the iris area. To that end, we use RANSAC [4] with three parameters (a, b, c) to fit a polynomial $ax^2 + bx + c$ to I_{up} and I_{low} . As we expect the upper eyelids to be concave and the lower to be convex and the curvature to be small we require that $a > 0$ and $b < 0$ for upper eyelids and $a < 0$ and $b > 0$ for lower eyelids and also $0.001 < |a| < 0.01$. At each iteration three points are sampled and a vector of residuals is returned for each pixel in the threshold edge image. When a few parallel edges for the eyelids are found, as in Fig. 2(b), the objective function for I_{up} penalizes estimations if a similarly shaped curve below them is found. Similar ideas are applied to I_{low} .

The pixels above and below the upper and lower eyelids respectively in I_{iris} are removed, in order to fit a circle to the remaining pixels. We use RANSAC with three parameters (x, y, r) and at each iteration sample three points. Similar ideas to the eyelids fitting are used, imposing restrictions over the points distribution and limiting the radius to the expected range of iris's radii.

After the iris has been detected, we use Levenberg-Marquardt (LM) [4] to do the fine tuning by iteratively minimizing a non-linear function for fitting a circle. The initial estimation (x_0, y_0, r_0) is the one found by the RANSAC. The main idea is to minimize the sum of distances of pixels with strong magnitudes to the circle (see more details in [13]). In order to choose as many pixels as possible that are actually on the iris, we add the following improvements:

1. The pixels with strongest gradients used for the LM minimization will not always fall only on the iris, as the eyelids and textures on the sclera have stronger edges. To that end, the y axis is scanned, and some weaker gradients are taken, if no other pixel was taken at that row.
2. Fittings that fall on pixels are preferred to interpolating between pixels. For instance, if the eyelids are not removed completely, there might be two strong edges at the bottom of the iris and an unwanted fitting will interpolate between the lines. It is partially solved by discrediting such residual vectors.

In Fig. 2(c), the pixels that were chosen for the fitting of the iris are shown in white. The input image has been processed to remove the eyelids and noise as described before. The final fitted curves to the iris and the eyelids are shown in Fig. 2(d). The detection of the iris in the rest of the sequence is done in a similar way, each time using the previous LM result as an initial guess. This results in the detected iris for the whole sequence. For every image, only the fast LM algorithm is used and not the relatively slower RANSAC we used for the detection of the eyelids.

2.2 Computing image translation

After locating the iris in each of the images, it is possible to align the images over the iris area. The main idea is to align the images by using the grey levels of the iris and eyelid areas to find the best homography between the images by minimizing differences of intensities in different scales. In order to prevent accumulating error, the images are taken in blocks, each

time aligning to the first image in the block. This step is very important as the camera is hand-held and the patient can move his gaze during the recording. In the aligned video, the iris is fixated in the same location and will have approximately the same radius at each frame. By doing so, it is possible to detect the BUT and the areas of break as explained in the next section. We refer the reader to [13] for more details.

2.3 Segmentation of the dry areas

An outline of the segmentation algorithm is given below and elaborated later:

Algorithm *Find Dry Areas*

Input: A video of n aligned images between two consecutive blinks.

Output: Dryness images of corneal break and black line. BUT.

1. Initialize BUT image: $\forall_{(x,y)} T(x,y) = n + 1$ (No break for all pixels).
2. Given the first image I_1 , find the region of interest (ROI) by finding the intersection between the eyelids and the iris.
3. Divide the ROI into 3 areas: upper eyelids, lower eyelids and corneal area, see areas numbered I-III in Fig. 3(a).
4. Calculate the average intensity for each of the three areas. (* To be used later for the creation of the dryness images *)
5. Repeat lines No. 2-4 for the last image I_n .
6. **for** each pixel $\in$ ROI
7. Calculate the three terms used to estimate the pixel's degree of thinning: $D(x,y)$, $B(x,y)$ and $R(x,y)$ (* See below *).
8. **for** each pixel $(x,y) \in$ ROI
9. **for** each image $i \leftarrow 1$ **to** n
10. intensity
 $\leftarrow \text{calcIntens}(D_1^i(x,y), B_1^i(x,y), R_1^i(x,y))$
11. **if** intensity $> B_T$ **and** $T(x,y) > i$
12. **then** $T(x,y) \leftarrow i$ (* Update the BUT for the pixel *)
13. Calculate the final dryness value for (x,y) .
14. Find black line in first and last images.
15. Find BUT and cause (corneal dryness, upper or lower eyelids) and statistics.

2.3.1 Producing the dryness images.

The dryness image is an image which shows the degree of dryness for each pixel in the iris area. An intensity value is calculated for each pixel, which is related to the degree of thinning. The higher the value, the drier it has become. In line 4, the average iris intensity for the first image, A_{I1} is computed. The brightest 1/8 pixels are omitted from the calculation, as errors in the alignment, the bright two circles in the iris area

(see Fig. 2(d)) and the inclusion of parts from the eyelids (very bright pixels) can bias the result. The choice of this value is based on approximating the ratio between the wanted and unwanted pixels in each area. The dryness image is given by:

$$\tilde{I}(x,y) = \lambda_1 D(x,y)(T_D/A_{I1}) + \lambda_2 B(x,y) + \lambda_3 R(x,y), \quad (1)$$

1. $D(x,y)$ is the average difference in intensities between the 4 last images and the 4 first images. However, as the average intensity for the iris can be as high as 150 or as low as 60, the term is made relative to the iris average intensity by multiplying by T_D/A_{I1} . The value of T_D has been set through learning to a value higher than the brightest iris average. A higher value for T_D , will result in a fast increase of $\tilde{I}(x,y)$ to images with low average iris intensity, A_{I1} . Therefore, such images are more sensitive to small changes in the pixels' intensity throughout the video.
2. The bonus map, $B(x,y)$ relates the pixel initial intensity to the iris initial average intensity. It is used to compensate for pixels whose initial value is lower or higher than the average value. For example, pixels with lower intensity than the average will usually become drier more quickly.
3. The downs-ups ratio $R(x,y)$ provides information how the pixel's intensity changes throughout the sequence (see [13]). Its contribution to $\tilde{I}(x,y)$ is usually low and it is mainly used for refinement. However, it is also used to detect or reject pixels near the borders whose intensity fluctuates for every small misalignment.

The value of B_T (line 11 in the algorithm), the minimum intensity for a pixel to be considered as a break, has been chosen arbitrary to be close to the maximum possible pixel intensity in the iris, and the calculations were adjusted according to it. An example for a dryness image is given in Fig. 3(c).

$\tilde{I}(x,y)$ is aimed at having a value higher than B_T when there is a break in (x,y) . To that end, the weighting coefficients have been adjusted by comparing to manual detections of the BUT by an optometrist. The parameters have been learnt by using a large set of training images and the coefficients were set according to the importance of each term. However, we believe that setting the parameters in a way that will match the clinician at all times is an ill-defined problem. This is due to the disparity between the BUT values given by different clinicians and sometimes even by the same clinician at different times. One reason is that the clinician tends to focus on a certain area and can miss break

areas that have a faster progress at other parts of the cornea. Another reason is that it is visually hard for the clinician to judge how black an area appears and to decide when a break happens. We are currently working on a new and more accurate definition to a BUT value calculated by a machine.

2.3.2 BUT computation.

The BUT computation is closely related to the way the dryness image is built. The intensity value in line 10 is calculated in the same way as for the final dryness image, however using only values from the first frame to the current frame (for example, $D_1^i(x, y)$). If the intensity of a pixel (x, y) is above B_T , the break-up threshold, the pixel is assigned (in line 12) the current frame i , meaning a break is suspected. However, if a pixel (x, y) shows further decrease in intensity at time $i > T(x, y)$, it is adjusted to have a later BUT, as areas devoid of fluorescein cannot have further breaks (or they are not expected to become any darker). Assuming S_{iris} is the total number of pixels in the iris, the BUT is the time passed until a small number of pixels are over the break threshold B_T :

$$BUT = \min_i \{ |T(x, y) \leq i| > S_{iris} T_S \} \quad (2)$$

The value of T_S is set to 0.0003, which is between 90 to 140 pixels in our images. This threshold is chosen to make sure that the BUT is computed correctly even when there are some misaligned pixels. This value can be correlated to the actual area that has the break. The diameter of the iris in over 90% of the population is between $9mm$ to $13mm$. The visible area of the iris in our videos (not covered by eyelids) is on average 88% of the iris area. Therefore, the minimum area that is required for a break ranges between $0.0168mm^2$ to $0.035mm^2$. This value makes sense, as smaller areas can be hardly detected by the clinician for a break.

In Figs. 3(c&g), the break areas are highlighted by dark grey on top of the dryness image. We divide the iris area into 5 areas in a similar way to the CCLRU standards (Contact Lens Research Unit) in [11]. The area of thinning and break in each quadrant is calculated and reported back to the clinician to be used in a follow-up inspection. Figs. 3(d&h) show graphs of the progress of the dry area as a function of the time passed since the blink. Each graph depicts both the area of thinning and the area of break, which is directly related to the computation of the BUT.

2.3.3 Black line detection.

The detection of the black line is performed separately because: (a) It clinically requires a different treatment

and (b) accurate detection of the black line is difficult due to the thin size and the proximity to the eyelids. The ground for the detection of the black line is prepared in lines 3-5, where the average values of the areas near the eyelids are calculated for I_1 and I_n . The areas searched for dryness are shown in Fig. 3(a), numbered by I and III. In most cases we expect to have quite a thin break line (if at all) after the blink and possibly a bigger area before the next blink. To that end, the search area for the black line after the blink is bounded by the iris and the two closest curves to the eyelids. The search area before the next blink is bounded by the iris, the eyelids and the furthest curve. We also use the same area to calculate the averages A_{U1} and A_{L1} , for the upper and lower parts respectively.

A pixel $I_1(x, y)$ near the upper eyelids is segmented if there is noticeable difference in intensity compared to the area average intensity:

$$7A_{U1}/8 > \frac{1}{4} \sum_{k=1}^4 I_k(x, y). \quad (3)$$

Therefore the average pixel intensity in the first 4 images has to be lower than $7/8$ of the area average (A_{U1}). We found out that since the averages A_{U1} and A_{L1} are computed in a small area, this constant holds well in most cases. Accurate detection of the eyelids (see Sec. 2.1) over the iris area is crucial in order not to include the bright pixels of the eyelids. The equation handles changes in intensity between images and between the upper and lower eyelids. It is also sensitive enough to discover thin break lines (see Fig. 4). Similar ideas are used for the lower eyelids and for both eyelids in I_n . A dryness value is associated with each pixel, and a break is assumed when it is higher than B_T , correlating the results to those of the BUT. We produce two images for I_1 and I_n , and by alternating between them, the clinician can see the changes in area and direction. Finally, if no corneal break occurred, but a break was detected near the eyelids, we report a break near the upper or lower eyelids, which probably caused the imminent blink.

3 Results

We captured over 100 videos of more than 30 patients. The patients were of different age groups and varied from having a very dry eye to no dryness at all. About 40 percent of the patients had dryness symptoms related to the black line. The patients were instructed to keep their eyes open as long as possible. A BUT of over 30 seconds has no clinical significance, so the captured videos are limited by 35 seconds. The




Figure 3. Dryness detection results of two different patients: (a&e) the image immediately after a blink. The three areas described in the algorithm in Line 3 are plotted. (b&f) The image before the next blink. (c&g) the dryness image. For clarity, the break areas are painted in dark grey. The other areas show different degrees of thinning (the brighter, the drier it has become). (d&h) Graph of the evolution of the dryness: the graphs show the computed area (as a percentage of the iris area) that has any degree of thinning and the area of the break, as a function of the time passed after the blink.

procedure was repeated on both eyes, and some patients were examined again at a later time, so the results could be compared. In [13], we compared the area of dryness to the hand segmented area and explained the objective problem in using such a measure. In this work, we compare the BUT value provided by the clinician to our computed BUT. We also asked the clinician to provide the reason for the break and compared it to our analysis. The possible reasons are: no break, a corneal break and a break originated near the upper or lower eyelids. A qualified clinician detected the frame that the break happened in 21 videos and wrote which area had the break (1 to 5). Breaks in areas 4 and 5 (see [11]) mean that it is a meniscus related break, so we compared it to our analysis of the origin of the break. The results are presented in Table 1. Analyzing the results show that the estimation of the BUT is in agreement with the manual detection. Taking into account that clinicians measure it in seconds, an error of 1 to 2 seconds is in the acceptable range. Also the origin of the break has been detected correctly in most videos. We believe that in some cases our results are more accurate than those of the clinicians, as they focus on a specific part of the cornea and have to scan to detect small changes near the eyelids (as in Fig. 3(c)). Moreover, our algorithm copes well with dry areas which form slowly, whilst it is hard for a clinician to detect the slow delicate changes.

It is worth mentioning that our algorithm provides more elaborated results related to the degree of thinning and analysis of the area and location, but it is hard to evaluate the correctness of the results objectively. We plan clinical trials to determine their objective value.

Table 1. Summary of results (1). The table compares the error in seconds in estimating the BUT. The average length of the videos is 12 seconds ranging from 6 to 20 seconds.

BUT estimation error	Patients number
< 1 second	14
1-2 seconds	4
2-3 seconds	2
3+ seconds	1
Total	21
Average Error	1.06 sec

4 Further Research

This paper presented an automatic method for analyzing different signs related to dry eyes: finding the BUT, the levels of thinning of the tear film, thinning near the reservoirs and analyzing the evolution of the dryness. We are currently extending our method to include more dryness signs related to corneal dryness and the tear reservoir. Then, in order to prove the method clinically useful, we aim to run clinical trials under the supervision of an ophthalmologist.

References

- [1] E. Bitton and V. Lovasik. Longitudinal analysis of precorneal tear film rupture patterns. *Adv Exp Med Biol.*, 438:381–389, 1998.
- [2] A. J. Bron, M. B. Abelson, E. Pearce, G. Ousler, A. Tomlinson, N. Yokoi, J. A. Smith, C. Begley, B. Caffery, K. Nichols, D. Schaumberg, and O. Schein. Methodologies to diagnose and monitor dry eye disease: Report of the diagnostic methodology subcommittee of the international dry eye workshop (2007). *The Ocular Surface*, 5(2):108–152, 2007.
- [3] J. Daugman. The importance of being random: statistical principles of iris recognition. *Pattern Recognition*, 36(2):279–291, 2003.
- [4] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [5] F. Holly. Surface chemical evaluation of artificial tears and their ingredients. iii dynamic properties. *Contact and Intraocular Lens Medical Journal*, 5(1):21–33, 1979.
- [6] F. J. Holly. Formation and rupture of the tear film. *Exp Eye Res.*, 15:515–525, 1973.
- [7] M. Lemp and J. Hamill. Factors affecting tear film break up in normal eyes. *Arch Ophthalmol*, 89(2):103–105, 1973.

Table 2. Summary of results (2). The table compares the reason for the break for patients without corneal break (11 out of 21). We compare the origin of the break (if any) and how many were detected correctly (A break can happen at both the upper and lower eyelids).

Region of break	Detected/Patients Number
None	5/5
Upper	5/6
Lower	1/2
Total	11/13




Figure 4. Results of detecting meniscus related dryness: (a) the image immediately after a blink. (b) The image before the next blink. (c) The black line after the blink. (d) The black line before the next blink. In this video the break is a result of thinning in the upper eyelids only (no corneal break). The visible increase of the area of the dryness near the upper eyelids is of clinical importance.

- [8] L. Ma, T. Tan, Y. Wang, and D. Zhang. Efficient iris recognition by characterizing key local variations. *IEEE Transactions on Image Processing*, 13(6):739–750, 2004.
- [9] M. Norn. Desiccation of the pre corneal film. *Acta Ophthalm.*, 47(4):865–880, 1969.
- [10] N. Ritter, R. Owens, P. van Saarloos, and J. Cooper. Location of the pupil-iris border in slit-lamp images of the cornea. In *ICIAP*, pages 740–745, 1999.
- [11] R. Terry, C. Schnider, and B. A. Holden. The CCLRU standards. *Optician*, 206(5430):18–23, 1993.
- [12] K. Tsubota. Tear dynamics and dry eye. *Progress in Retinal and Eye Research*, 17:556–596, 1998.
- [13] T. Yedidya, R. Hartley, J. Guillon, and Y. Kana-gasingam. Automatic dry eye detection. In *MICCAI (1)*, pages 792–799, 2007.