

AUTOMATIC QUANTIFICATION OF PUPIL DILATION UNDER STRESS

Julien Jomier, Erwann Rault and Stephen R. Aylward

Computer-Aided Diagnosis and Display Lab, Department of Radiology
The University of North Carolina at Chapel Hill
27599 Chapel Hill, USA

ABSTRACT

Response of the autonomic nervous system can be assessed by quantifying the dilation of the pupil of a human placed in darkness. When the intensity of the surrounding light is low, human's pupil dilates to capture a maximum number of photons. This type of dilation is known as dark adaptation. The aim of this study is to test the hypothesis that dark adaptation is slowed proportional to the amount of stress that an individual has experienced. We have developed a new tool for automatic quantification of pupil dilation. Our system uses statistics of the color of the pupil and the iris as well as model-to-image registration and region growing segmentation techniques. Our system has been tested on several images and shows high correlation with hand segmentation.

1. INTRODUCTION

The National Vietnam Veterans Longitudinal Study seeks to quantify the relationship between stress and the long-term degradation of the autonomic nervous system. This multi-central study will conduct many autonomic nervous system measures; we are particularly interested in testing the hypothesis that dark adaptation is slowed proportional to the amount of stress that an individual has experienced. We seek to automate the measurement of pupil and iris areas from color digital photos so as to calculate the ratio of those areas as a measure of the amount that an individual has dark-adapted.

Individuals are placed in a darkened hood, Fig.1-left, and after 30 seconds a color digital photo is acquired to measure their amount of dark-adaptation. The camera is tilted to approximately standardize the vertical location of the eyes in the photo. A ring flash is used to minimize the specular highlights within the photo and to limit generally those highlights to the center of the pupil. The photos are 4256x2848 images with 24 bit RGB pixels and are stored using jpeg compression to reduce file size from 35.5 megabytes to 4.3 megabytes. For the work presented in this paper, 5 images were used for training and 15 for testing.



Fig. 1. *Left:* Acquisition Setup *Right:* Original image from the Fuji S1 camera

2. METHODS

This section describes each step of the algorithm, from eye localization to precise pupil and iris segmentation. First the pupil is coarsely segmented using a model-to-image registration approach then statistics on the RGB components of the pupil and iris are refined and detailed segmentation is performed using active contours on the probability maps. The same algorithm is used to segment the iris using the color probability density maps of the iris.

2.1. Eye localization

The first step of the algorithm is to locate eyes within the image. The image is first sub-sampled by a factor of 5 to decrease computation time. The color characteristics of pupil are used to determine the position of the two pupils. In all images of our training set we have found that the red component is always less or equal to the blue component inside the pupil. Therefore, taking all pixels such that $Red - Blue < \rho$ results to a rough segmentation of the pupil, Fig.2-left. The threshold parameters ρ is determined during the training stage. Our acquisition setup ensures that the level of the pupils is roughly fixed and that the left eye will be in the left part of the image and the right eye in the right part correspondingly, so we can easily separate the two eyes. A center of mass computation is performed on the two sides of the binary image obtained previously. This step defines the center of our regions of interest. The size of these regions remains fixed. After this stage, the two eyes

are treated independently of one another.

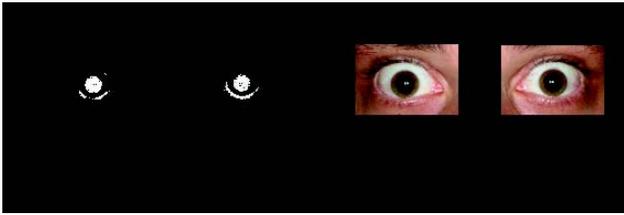


Fig. 2. *Left:* Binary image, *Right:* The two regions of interest obtained after eyes selection

2.2. Coarse Pupil Segmentation

During this stage we segment the pupil. We first create a binary image of the eye to segment the pupil. We use a different threshold than the previous step. Here, we are setting to one all pixels for which $\frac{(Blue+Red)-\mu}{\sigma} > 0$, and set to zero otherwise. In most of the cases, the resulting binary image, Fig.4-Top-Right, contains holes and is not a perfect circle. Therefore, we perform holes filling using morphological operators, Fig.4-Bottom-Left. We then use the Hough transform [1] to find the center of the pupil. However, because the circle is not well defined and the radius as well as the position of the center are unknown, the resulting parametrization is not perfect. Therefore, we use a model-to-image registration technique to refine the parametrization. This technique is described next.

Our model-to-image registration uses a circle as a model. The parameters to be optimized are the position of the center $X = (x, y)$ and the radius r . We use the following metric to quantify the fit of the circle to the binary image.

$$f(r, X) = \frac{A(r + \Delta r, X) - A(r, X)}{2 \cdot r + \Delta r} \quad (1)$$

where $A(r, X)$ is the number of the bright pixels inside the circle of radius r centered at X .

Therefore the optimization process is minimizing the derivative of f with respect to r and x . The computation of the derivative can be done using finite differences, Fig.3 shows the plots of $f(r)$ for a fixed X . The top plot represents $f(r)$ with $0 < r < 6$ when the optimal radius R is equal to 3. This plot shows that f is maximum inside the circle to be found and drops to zero outside. Therefore by taking the derivative of f , Fig.3-Bottom, we show that our metric is minimal at the edge. These two plots illustrate the perfect case where the circle is perfectly filled in and no noise is present in the background. However, our metric shows accurate and robust results on real images. The result of the optimization process is presented in Fig.4-Bottom-Right. In our experiments, we use a one-plus-one

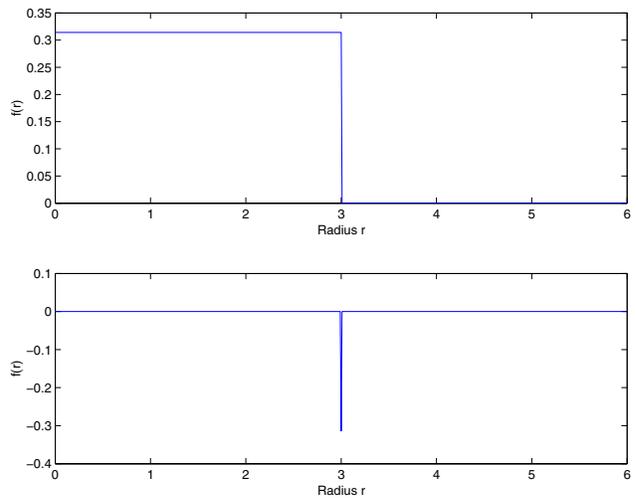


Fig. 3. *Top:* Metric values of f for a radius range from 0 to 6 where the circle to be found has a fixed radius of 3 *Bottom:* Derivative of the metric obtained using finite differences. The derivative of f is minimum for the optimal radius ($r = 3$)

evolutionary optimizer [2]. After each iteration of the optimization process we recompute the mean and the standard deviation of the pixels within the estimated circle to update. To ensure that samples belong to the pupil class we perform a uniform sampling on a circle with a radius equal to half of the estimated radius. We also discard pixels that have high RGB components because these pixels are more likely to belong to the highlight. Therefore a new binary image is created by thresholding the probability density map.

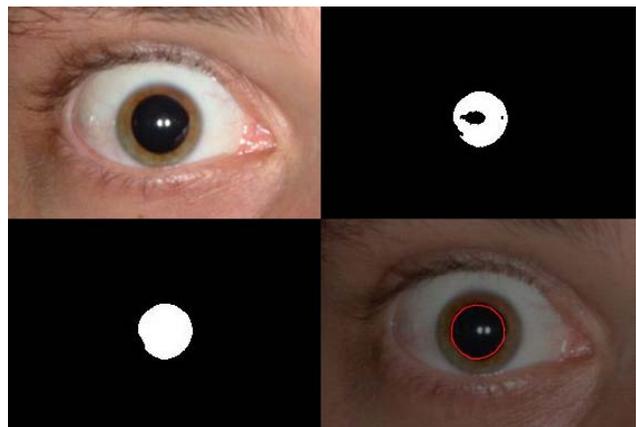


Fig. 4. *Top-Left:* Region of interest image, *Top-Right:* Binary image obtained using our special RGB threshold, *Bottom-Left:* Binary image after closing *Bottom-Right:* Resulting circle registered with the pupil edge

2.3. Iris Segmentation

Once the pupil is segmented, we can estimate color statistics of the iris by looking at the pixels in the region around the pupil. However, since some subjects are not able to completely open their eyes, the top or the bottom part of the iris may be hidden. Therefore we only probe the pixels on the left and right regions of the pupil. The beginning of the region is defined by a circle of radius equal to the radius of the pupil plus 5 pixels to avoid taking pixels which belong to the pupil. The end of the region is defined by another circle whose radius R is determined via another optimization process. Fig.5-Top-Right shows the two regions used to sample the iris. We use, once again, the circle-to-image registration method described in 2.2. During this optimization, however, the center of the circle is kept fixed as the center of the pupil and only the radius R is optimized. Again, the statistics of the iris RGB pixel values are re-estimated at each iteration during the registration. Fig.5-Bottom-Right shows the result of the registration.

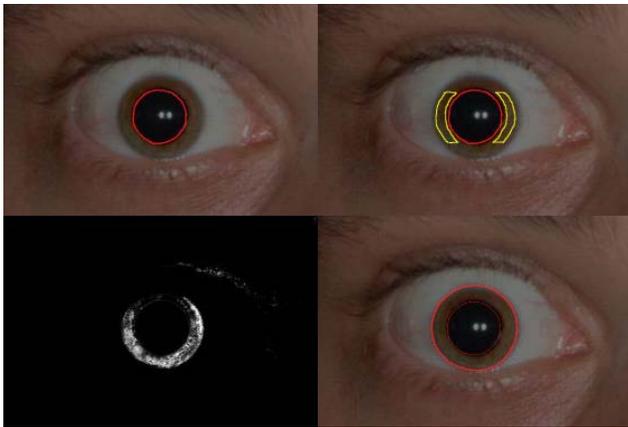


Fig. 5. *Top-Left:* Circle matching the pupil after model-to-image registration, *Top-Right:* Regions used to compute the probabilities of the iris, *Bottom-Left:* Probability density function of the iris, *Bottom-Right:* Resulting circle registered with the iris

2.4. Precise Pupil Segmentation

At this stage we have two circles which define the pupil and the iris and the probability density functions of these two entities are well defined. The goal is to compute an optimal surface of the pupil using an active contour. We use, in this part and for the rest of the processing, the original image to take advantage of the full resolution of the captured images. One can notice that we are not refining circles on the original image because we found that the coarse segmentation was accurate enough. We seek the boundary between the iris and the pupil probability maps, therefore we form

a difference image between the pupil probability map and the iris probability map, Fig.6-Top-Right. In most of the images, the boundary between the pupil and the iris is not defined by a sharp edge. We threshold the probability map to form a binary image so that every pixel with a probability higher or equal to zero are set to one and zero otherwise. In order to remove the highlight and the possible noise in the binary image we perform morphological operations. The resulting image is shown in Fig.6-Bottom-Left. This binary image is then use as a gradient map by the geodesic active contour [3] [4]. The active contour is initialized by an implicit surface with a radius of $\frac{3}{4}$ times the initial radius of the pupil. Fig.6-Bottom-Right shows the final result of the segmentation. Since the iris may be partially obscured by a lid, we simply use the result of the model-to-image registration process. As a final step we compute the ratio of the number of the pixel in the iris and the number of pixel in the pupil.

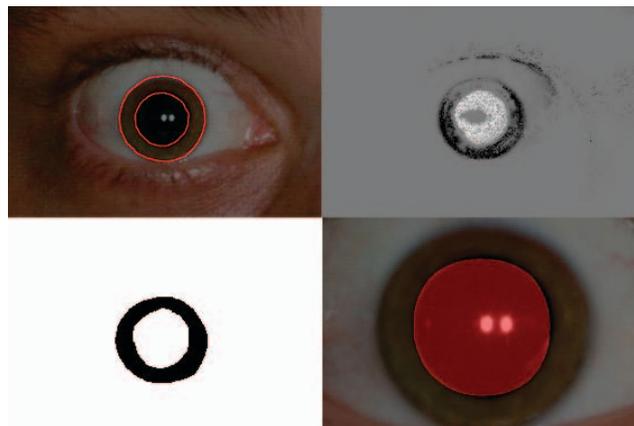


Fig. 6. *Top-Left:* Circles matching the iris and the pupil after model-to-image registration, *Top-Right:* Image representing the difference between the probability map of the iris and the pupil, *Bottom-Left:* Binary probability map after morphological operations, *Bottom-Right:* Segmented pupil via active contour

3. RESULTS

We trained our algorithm on 5 images using only the left eye and tests were performed using 10 images, 20 eyes, including the training. We compare our algorithm to hand-segmentations of 5 raters. Results are shown in Figure 7.

For each eye, we consider the mean area of the hand-segmented pupil as the ground truth. Each hand segmentation is compared to the ground truth to evaluate the relative error. The average error over all raters as well as the maximum error show a high accuracy among the raters. Our method has a slightly higher error rate but a lower maxi-

Segmentation	Average μ	μ Error	Max Error
Raters	35010	1.94%	6.76%
Computer	34780	2.41%	5.80%

Fig. 7. Results obtained on 10 subjects (20 eyes). The subjects represent a large variation of iris and pupil colors. The mean average is expressed in pixels and represents the average of all segmented areas of the pupil. The mean and the max errors represent the relative error of the segmentation compared with the ground truth.

num error. This can be explained by the smoothness constraint used in the active contour algorithm as well as the morphological operators used to create the edge map that tend to smooth the boundary of the segmentation. On the other hand, specular highlight produced by the flash can affect the hand segmentation. Moreover, for some individuals, pupil and iris do not define a sharp boundary and the color of the pupil tends to diffuse non uniformly in the iris, making the segmentation difficult. One can notice that the error of the segmentation is located on the boundary of the circular object and therefore is proportional to the radius. Our experiment has shown that the area of the pupil can differ between two individuals by a factor of two. Therefore our error measurement is evaluated as a relative error. One strength of the system is the segmentation time. It takes approximately 20 minutes for a rater to hand trace the 10 eyes. Our algorithm computes the area of the pupil and the diameter of the iris in less than one minute per eye with the same accuracy. Moreover, our tool is fully automatic making it useful for processing hundred of images without human intervention.

4. DISCUSSION

We have implemented a new tool for automatic pupil dilation quantification. Our system uses a model-to-image registration metric combined with statistics and level-set segmentation. Our model-to-image metric shows robustness and accuracy for all the images. We are conscious that the validation is weak at this moment. We are continuing to capture and process images. This tool will be used throughout the United States to quantify dilation of the pupil for Vietnam veterans leading to a full validation of the hypothesis that dark adaptation is slowed proportional to the amount of stress that an individual has experienced.

This work was developed using the Insight Toolkit [5].

5. REFERENCES

[1] D. H. Ballard, "Generalized hough transform to detect arbitrary patterns," *IEEE Transactions on Pattern*

Analysis and Machine Intelligence, vol. 13, no. 2, pp. 111–122, 1981.

- [2] M Styner and G. Gerig, "Evaluation of 2d/3d bias correction with 1+1es-optimization," *Technical Report BIWI-TR-179*.
- [3] Ross T. Whitaker, "Algorithms for implicit deformable models," in *Fifth International Conference on Computer Vision*. IEEE, 1995, IEEE Computer Society Press.
- [4] Ross T. Whitaker, "A level-set approach to 3D reconstruction from range data," *International Journal of Computer Vision*, vol. 29, no. 3, pp. 203–231, 1998.
- [5] Insight Software Consortium, "The insight toolkit - segmentation and registration toolkit," www.itk.org.