Graefe's Arch Clin Exp Ophthalmol (1993) 231:90–94

Graefe's Archive
for Clinical and Experimental
Ophthalmology

Automated detection and quantification of retinal exudates

Russell Phillips[1], **John Forrester**[1], **Peter Sharp**[2]

[1] Department of Ophthalmology, Medical School, University of Aberdeen, Foresterhill, Aberdeen, AB9 2ZD, Scotland
[2] Department of Biomedical Physics, University of Aberdeen, Foresterhill, Aberdeen, AB9 2ZD Scotland

Received: 19 December 1991 / Accepted: 21 September 1992

Abstract. Retinal exudates are a common manifestation of vascular damage in a variety of retinal diseases. We have used computerized image analysis to detect and measure the area of exudates from digitized colour fundus slides of patients with diabetic retionpathy and have assessed the repeatability, reproducibility, and accuracy of the technique. The analysis was entirely independent of the operator apart from choice of the region to be analysed. The coefficient of variation for repeatability was between 3% for large areas of exudate and 17% for small areas of exudate. The reproducibility was also within this range. Sensitivity was between 61 and 100% (mean 87%). False-positives were observed in 5 of 30 regions analysed, and these could have been eliminated by using more stringent criteria for selection of images for analysis. Time taken for the analysis was approximately 3 min.

Introduction

Hard retinal exudates are a characteristic feature of many retinal disorders, including diabetic retinopathy. In general, they are associated with patches of vascular damage with leakage, especially around microaneurysms. Consequently, their size and distribution may vary during the course of a retinal disorder or in response to therapy designed to limit vascular leakage, e.g. focal laser therapy. However, accurate quantification of the degree of retinal exudation has limited precision and accuracy with existing clinical techniques [1]. In this study digital imaging techniques have been applied to estimate the area of retinal exudates in a fundus image. Whereas previous papers on computerized exudate analysis [3, 7] have not rigorously assessed the clinical usefulness of the techniques, we have made this a major part of the current study.

Two generalizations can be made about exudates. Firstly, large confluent areas of exudate are usually of a very much higher grey level than most of the rest of the image and can often be easily segmented using a high threshold, the same value being applied to the whole region of interest. Secondly, the smaller exudates usually have a lower grey level. However, although they are of much lower grey level than the majority of large confluent exudates, they are of higher grey level than their immediate surroundings. Thus by using threshold values peculiar to subregions within the full region of interest it is possible to segment them. These generalizations are confounded by two important image characteristics. Firstly, there may be present in the image large areas of confluent exudate that are of relatively low intensity. Secondly, the image might contain very marked pigment epithelial mottling or very marked non-uniformity of illumination. In the first case, the exudates are detectable with a high degree of sensitivity by using lower thresholds, but this approach leads to loss of specificity in those images with non-uniformity. Conversely, if the thresholds are raised to maintain specificity in non-uniform images, sensitivity is lost in those with large low-intensity exudates. The program described in this paper is biased towards specificity at the cost of sensitivity. Possibilities for improving sensitivity without losing specificity will be discussed later.

Materials and methods

Acquisition

Colour slides of a 35° or 50° field centred on the macula were projected through an Optilas type 48 red-free filter and digitized using a modified version of a system described previously [4, 6], based on an IBM compatible computer fitted with a frame grabber, the image being captured using a CCD camera. The camera exposure was adjusted so that the maximum grey level of any feature

This work was supported by grants from: The Scottisch Home and Health Department, The Scottish Hospital Endowments Research Trust, The TFC Frost Trust, and Grampian Health Board.

Correspondence to: J. Forrester

within the image (i.e. exudates or optic disc) was in the range 234 to 237 (the maximum possible grey level value is 255, a higher intensity causing saturation of the digital image). In normal images (no exudates) where the optic disc was pink, with no area of central pallor, the camera exposure was adjusted so that the mode of a well-centred region of interest 416×288 pixels was 75. The mode was felt to give the truest representation of the "background" grey level of the fundus image.

Data processing

Retinal exudates appear "whiter", i.e. have a higher grey level value, in the digital image than the surrounding normal retina. This is demonstrated in the profiles across a large dense exudate and small, less dense exudates shown in Fig. 1. While Fig. 1 demonstrates that the exudates have a higher grey level than their surroundings, the absolute value of the grey level is mainly dependent on two factors:

1. The vignetting of the projection system, which tends to make exudates nearer the centre of the field appear brighter.
2. The size of the exudate; larger exudates tend to have greater depth and density than smaller ones and therefore appear brighter.

This means that a simple grey-level thresholding, where exudates are identified as all areas in the image having a grey level exceeding a threshold value [5], would be unlikely to detect all the exudates within the region of interest. The program was therefore designed to incorporate a number of adaptations that allowed the computer to vary threshold according to the characteristics of the individual image. The solution presented represents a compromise between sensitivity and specificity because of the strict aim of producing a program that does not involve user interaction in the thresholding stages.

Preprocessing

Shade correction to eliminate nonuniformities of illumination was based on the method described by Jagoe et al. [2]. Firstly, the image detail was sharpened by convolution using the kernel:

$$\begin{array}{rrrrr} -1, & -1, & -1, & -1, & -1 \\ -1, & 1, & -2, & 1, & -1 \\ -1, & 2, & 18, & 2, & -1 \\ -1, & 1, & -2, & 1, & -1 \\ -1, & -1, & -1, & -1, & -1 \end{array}$$

The sharpened version of the original 512×512 image was reduced to 128×128 pixels by averaging blocks of 4×4 pixels and heavily smoothed using a 9×9 median filter [5]. This smoothed image was rezoomed to its original 512×512 size before being smoothed by convolution with an 11×11 mean filter [5]. The entire smoothing procedure was repeated on the smoothed image, producing even heavier smoothing. To produce a shade-corrected image the original sharpened image is divided by the smoothed image and normalized to a mean grey level of 100.

Unfortunately, this technique has the effect of reducing the intensity of large confluent areas of exudate. An adaptation was therefore made in which pixels in the smoothed image with a grey level value greater than (mode + 55) (i.e. only large confluent exudates, optic disc central pallor) were reset to (mode + 40). The smoothing process was then repeated twice.

Two types of shade-corrected image were produced: the first by dividing the sharpened image by the smoothed image and the second by subtracting the smoothed image from the sharpened image. The former image retains good detail and contrast for small exudates, but slight loss of contrast of larger exudates. The latter image is a generally low intensity image with much of the background image being at or near a grey level of 0. Exudates maintain a high contrast relative to the background. These two shade-corrected images were added together to enhance the contrast of the exudates in the divided image. In order to reduce the enhancement of relatively high grey level background pixels caused by this addition, the subtracted image was first offset by a constant of −25.

An example of the effect of this preprocessing is shown in Fig. 2a.

Thresholding

A region of 256×192 pixels was positioned by the operator over the area of interest in the image. Large confluent exudates were detected by thresholding using a threshold value of (mode + 50) for all the region. The thresholded pixels were stored in an empty frame buffer.

To detect the smaller, lower intensity exudates, local thresholding was used. The region of interest was divided into blocks of 32×32 pixels and then thresholded with a grey level value of (mode + 20), where the mode was that of the pixels within this small block. Thus the threshold might vary from block to block. These thresholded pixels were then combined with those previously stored, and the combined image represented what had been identified as the areas of exudates.

Display

The area of thresholded pixels representing exudates was calculated, and the thresholded pixels were overlaid on a copy of the sharpened image (Fig. 2b).

Assessment of performance

Repeatability was tested by acquiring and analysing 20 regions of interest in 14 different colour slides on five separate occasions. On the first acquisition of each different slide, the image was stored on disc and the coordinates of the regions analysed were noted. Subsequent acquisitions of the same slide were aligned with the first by overlaying the new on the old using software to shift the images until they were visually judged to be in alignment. This ensured that the same regions on the slide were analysed at each repetition. The standard deviation and coefficient of variation of the technique were calculated for each region analysed (Figs. 3, 4).

Reproducibility was tested by carrying out the acquisition and analysis on two slides taken 4 days apart of the same eye of a single patient (Table 1). The accuracy of the technique in terms of its sensitivity and specificity was difficult to test for three reasons. The first was the subjectivity of deciding whether some of the low-intensity exudates were true exudates or not. The second was deciding where some of the fuzzy-edged exudates ended. The third problem was that of accurately measuring the true area of exudates. An attempt was made to assess sensitivity and specificity on 30 regions of 13 slides. The computer analysis was performed; false-positive and false-negative exudates were identified and marked on a clear acetate sheet. This sheet was then overlaid on a sheet of graph paper, an estimate of the area of false-positives and negaives made (Table 2), and sensitivity calculated for each region (see Fig. 5).

Results

The results of analysing an image are illustrated in Fig. 2b, the time to perform this analysis on our system being about 3 min.

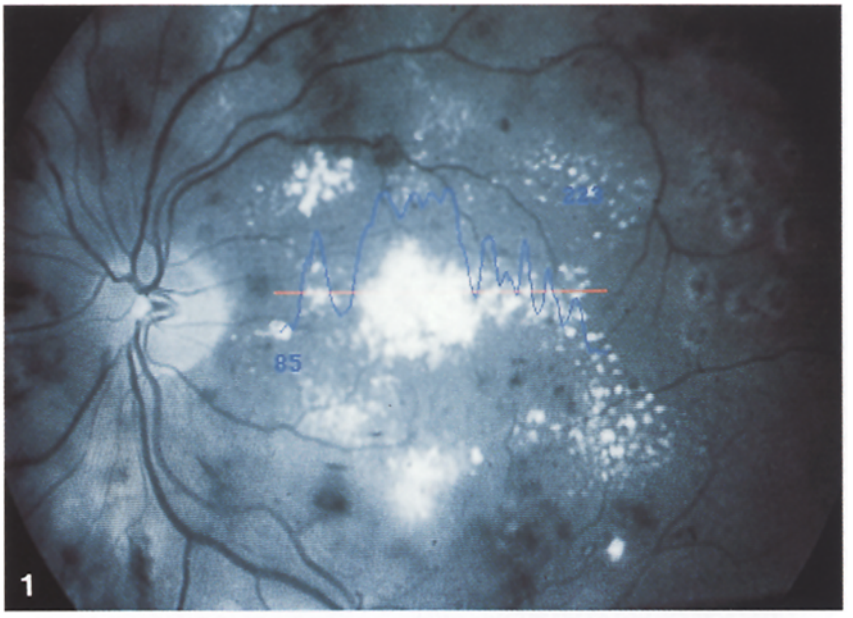

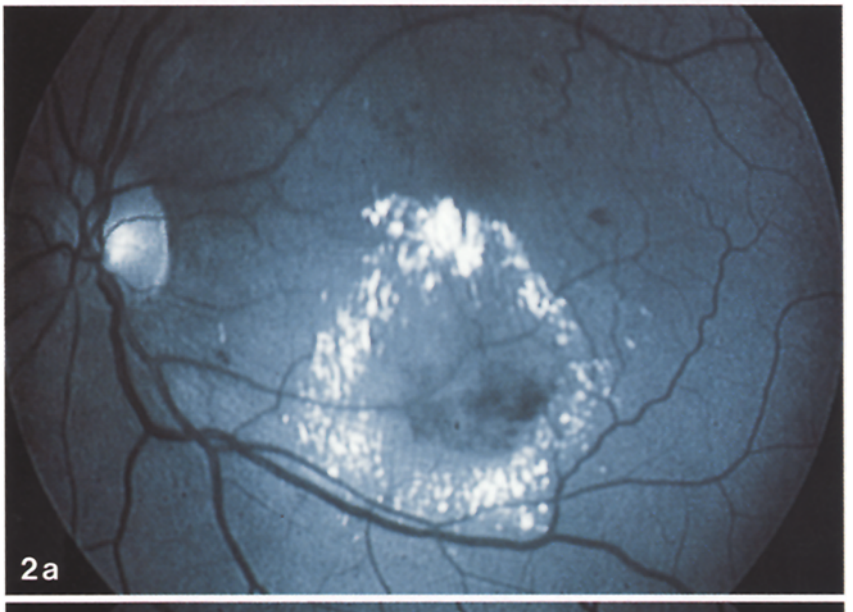

Fig. 1. Digitised colour slide showing retinal exudates. A grey-level profile along the *red line* displays the grey level values shown on the *blue line*. Note the variation in the actual grey level values of different exudates

Fig. 2. a Digitized colour slide showing exudates after "sharpening" but prior to processing for exudate analysis. **b** The result of exudate analysis

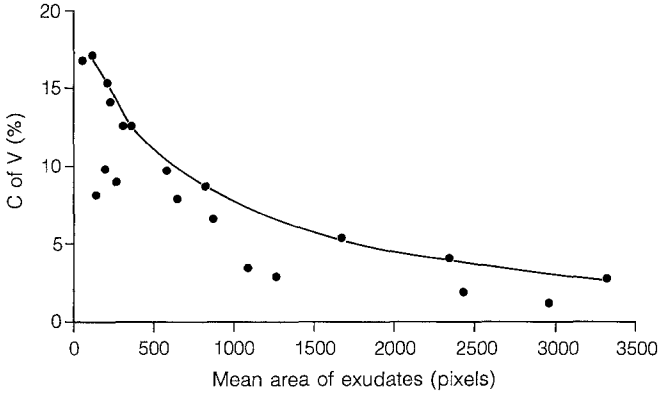

Fig. 3. Graph of coefficient of variation against mean area of exudates. Each *closed circle* represents a single region of interest that has been acquired and analysed on five different occasions. The *line* has been drawn through the highest points to give a "maximum likely" coefficient of variation attributable to the acquisition stage of the technique.

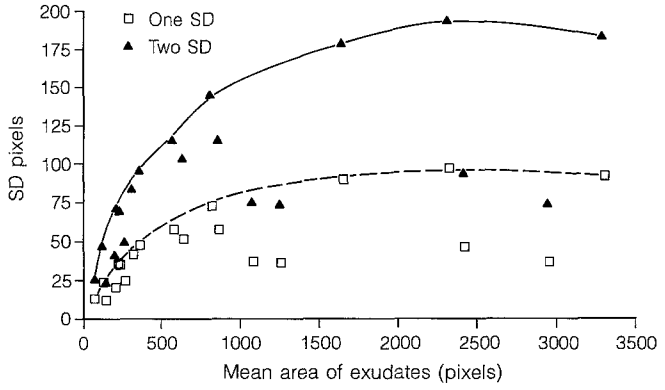

Fig. 4. Graph of standard deviation against mean area of exudates. Each *closed circle* represents a single region of interest that has been acquired and analysed on five different occasions. The *line* has been drawn through the highest points to give a "maximum likely" standard deviation attributable to the acquisition stage of the technique

Table 1. Reproducibility study

	Area of exudates in pixels	SD	C of V (%)
Time 0	155	4.4	2.8
Time 4 days	148	11.8	7.9

The results of the reproducibility studies described above are illustrated by the graph of coefficient of variation (standard deviation/mean) plotted against mean area of exudates (Fig. 3). In the case of exudate areas below approximately 100 pixels, while exudates were reliably and consistently detected, small variations in the numerical value of the area detected led to a larger coefficient of variation. Note that the curve has been plotted on the maximum coefficient of variation, rather than fitted through the data points. For the two slides of the same patient taken within 4 days of each other, the measured area of exudates was well within the limits of variability expected for this technique (Table 1), indicating that minor variations in photographic technique did not have a significant effect on the analysis of exudate area for this image pair.

Table 2. Assessment of performance of exudate program.

Area of exudates (pixels)	False-positives (pixels)	False-negative (pixels)	Sensitivity (%)
379	0	50	88
897	0	100	89
1138	0	100	91
1016	0	50	96
1191	0	50	96
1725	30$	50	97
207	30$	0	100
4153	0	0	100
574	0	50	91
786	0	100	86
490	60&	0	100
435	0	50	89
445	0	50	89
2032	0	350	85
2152	0	250	89
929	400#	100	84
1041	400#	100	86
507	0	100	83
631	0	100	86
7591	0	3000	71*
6201	0	3500	63*
5978	0	2000	74*
4647	0	2000	69*
8925	0	200	97
7969	0	200	97
8816	0	200	97
33	0	0	100
291	0	0	100
304	0	80	79
160	0	100	61
		Mean	87
		SD	11

An explanation of the results marked with the symbols $, &, # and * is given in the Discussion

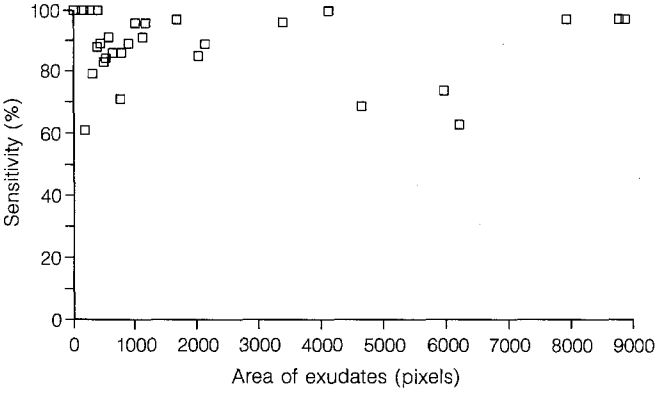

Fig. 5. Each data point represents a single region of interest. The lowest sensitivities (highest number of false-negatives) can be explained by the characteristics of the individual images

Discussion

An automated method for detecting and quantifying retinal exudates that requires no user interaction for the

thresholding stages has been developed. This represents a considerable improvement on previously described methods [7] where the necessity for user interaction in the thresholding stages reduces objectivity. The system proved to be reasonably robust, in that it could deal with slides of varying quality, and repeatable, with coefficients of variation for the acquisition process ranging from 17% for small areas of exudate (less than one-tenth disc area) to under 5% for large areas of exudate (exceeding one disc area).

False-positives were seen in five regions. The most significant numbers (marked # in Table 2) were found in an image where there was an area of dense retinal gliosis with a high grey level within the region of interest. It would be reasonable to exclude images such as this, or at least to modify the results. The false-positives marked $ in Table 2 resulted from the region of interest overlapping an area of photographic artefact on the slide. The false-positives marked & in Table 2 were the result of very prominent nerve fibre layer striations.

This system also detects other lesions with a high grey level intensity such as drusen, and areas of bare sclera such as peripapillary scleral crescents, or areas of gross retinal pigment epithelial/chorioretinal atrophy. Cotton wool spots are generally not detected because their grey level falls below the detection threshold. Mild-to-moderate depigmentation due to retinal pigment epithelial loss is not detected for the same reason.

Sensitivity averaged 87%, ranging from 61 to 100%. Unlike repeatability, sensitivity was not related to the exudate area (Figs. 3–5). False-negatives were the result of exudates of low grey level intensity, where contrast between the exudate and its surroundings was poor. It can be seen from Fig. 1 that some exudates are barely above the background grey level in the brighter parts of the image. The preprocessing stages of the exudate program attempt to maximize the information contained within the image. Four of the five worst results for sensitivity (marked with an asterisk in Table 2) were obtained from different regions of the same image, in which a large proportion of the exudates apparent in the photograph were of low contrast with the background. This was due, at least in part, to the exudates being located more peripherally in the photographic field. In spite of the low contrast exudates in this image, during program devolopment it was possible to obtain excellent sensitivity for this image, but the thresholds used led to an unacceptable level of false-positives on other images. The final program is therefore a compromise between sensitivity and specificity, biased towards specificity. These adverse results emphasize the necessity for some ophthalmological expertise in interpretation, and for strict criteria for selection of slides for analysis.

A number of modifications are now under active consideration to improve the performance of the system. Acquisition of the images direct from the patient would be advantageous in several ways. Firstly, by excluding the stages of photographic processing, projection, and reacquisition repeatability should be enhanced. Secondly, several digital images could be acquired and analysed at each photographic session, and the results of more than one image analysis averaged. In a previous study [6] this approach approximately halved the coefficient of variation for repeatability. Direct acquisition also allows the images to be assessed by eye for quality, and poor images can be rejected. The optics of the fundus camera and the eye produce some vignetting of illumination, which is manifest in photographic slides. It is possible to reduce this vignetting by placing a diffuser in the illumination pathway of the fundus camera, and this is being pursued. The scanning laser ophthalmoscope [8] is beginning to find its way into clinical use as a retinal imaging device. It is capable of producing images of high resolution, and its narrow band illumination and confocality could be used to great advantage.

The analysis part of the process could perhaps be improved using more complex thresholding techniques, but such methods are very intensive computationally and take an unacceptably long time on a modest computer. Local thresholding using a variety of subregion sizes was tried and found not to improve significantly on the results presented above.

The purpose of testing the system for repeatability of the acquisition process and reproducibility is to ascertain the limits of experimental variation, so that in the clinical situation the significance of a measured change in exudate area is known. On the basis of these experiments it has been possible to derive a measure of the expected experimental variation in exudate analysis, and to produce curves of coefficient of variation and standard deviation plotted against area of exudates that could be used in a clinical setting. Perhaps of more importance, these results represent a baseline against which to measure future improvements in both the acquisition and analysis sides of the process. The speed and objectivity of the present method for detecting and quantifying retinal exudates are significant advances in the assessment of diabetic exudative retinopathy.

References

1. The Diabetic Retinopathy Study Research Group (1981) A modification of the Airlie House Classification of diabetic retinopathy: Diabetic Retinopathy Study: report no. 7. Invest Ophthalmol Vis Sci 21:210–226
2. Jagoe JR, Blauth CI, Smith PL, Arnold JV, Taylor RM, Wooton R (1990) Quantification of retinal damage during cardiopulmonary bypass: comparison of computer and human assessment. Procs IEEE 137:170–175
3. Katz N, Goldbaum M, Nelson M, Chaudhuri S (1988) An image processing system for automatic retinal diagnosis. Procs SPIE 902:131–137
4. Phillips RP, Ross PGB, Tyszka M, Sharp PF, Forrester JV (1991) Detection and quantification of hypofluorescent leakage by computer analysis of fundus fluorescein angiograms. Graefe's Arch Clin Exp Ophthalmol 229:329–335
5. Rosenfeld A, Kak AC (1982) Digital picture processing. Academic Press, Orlando, pp 61–73; 254–263
6. Spencer T, Phillips RP, Sharp PF, Forrester JV (1992) Automated detection and quantification of microaneurysms in fluorescein angiograms. Graefe's Arch Cli Exp Ophthalmol 230:36–41
7. Ward NP, Tomlinson S, Taylor CJ (1989) Image analysis of fundus photographs. The detection and measurement of exudates associated with diabetic retinopathy. Ophthalmology 96:80–86
8. Webb RH, Hughes GW, Delori FC (1987) Confocal scanning laser ophthalmoscope Appl Optics 26:1492–1499