Feature-Based Registration of Retinal Images

ELI PELI, REED A. AUGLIERE, AND GEORGE T. TIMBERLAKE

Abstract—**Registration of retinal images taken at different times frequently is required to measure changes caused by disease or to document retinal location of visual stimuli. Cross-correlation has been used previously for such registration, but it is computationally intensive. We have modified a faster algorithm, sequential similarity detection (SSD), to use only the portion of the template that contains retinal vessels. When compared to standard SSD and cross-correlation, this modification improves the reliability of detection for a variety of retinal imaging modalities. The improved reliability enables implementation of a two-stage registration strategy that further decreases the amount of computation and increases the speed of registration.**

I. Introduction

OPHTHALMIC practice is second only to radiology in its reliance upon imaging science. However, the range of imaging applications in ophthalmology goes beyond X-ray, NMR, and ultrasound, because the transparent nature of the eye allows imaging in the visible spectrum.

The retina or fundus of the eye is photographed most commonly in clinical practice. White-light color photographs of the retina are used for general evaluation and as stereo images to allow measurements of depth. Red-free light illumination is used in photography of the nerve fiber layer and in fluorescein angiography with fluorescein injected intravenously. The wide availability and applicability of retinal images have led to the development of computer image-analysis techniques [1], [2], as well as a number of commercial image-processing systems for ophthalmology [3], [4].

Registration (alignment of retinal images taken at different times) is required frequently in image-processing applications. Images may be used to detect slow changes in the retina and thus may be separated by a few years. Changes such as these occur in drusen [5], nerve fiber layer damage [6], optic disk cupping and pallor [7], and intraocular tumors [8]. Registration also may be required for pictures taken only a fraction of a second apart, as in measuring dilution curves in fluorescein angiography [1], or at video rates as required in fundus perimetry using the scanning laser ophthalmoscope (SLO) [9]. Registration also may be used for frame averaging to improve signal-to-noise ratio.

Most researchers have used cross-correlation for fundus image registration [1], [10]. With this technique, the template array is moved across the search area, and the coefficient of cross-correlation is calculated for every point in the search area. The peak of the cross-correlation surface indicates the point of maximum correlation between the template and the search area. This peak often is relatively flat, so that preprocessing of the two images with high-pass filtering usually is required to obtain a sharp peak [1], [10], [11]. The cross-correlation technique is computationally intensive and may require a lengthy calculation period despite the use of fast mainframe computers [1], [12]. For this reason, various attempts to speed the calculation have been implemented. Parker *et al.* [13] used the fast Fourier transform to calculate the correlation via the frequency domain; Glazer *et al.* [10] suggested a parallel hierarchical correlation that could be implemented using parallel processing. However, they only simulated the technique on a serial computer and thus did not significantly reduce the calculation time.

Because of the lengthy calculations required, a number of researchers have used various manual techniques to register fundus images. In these techniques, one image is displayed in one plane of the image processor, while another is digitized continuously and moved manually until registration is achieved. Various display strategies can be incorporated to facilitate such registration. Algazi *et al.* [2] used a sequential monocular display technique in which the two images are flickered sequentially, and the flicker is reduced when proper registration is achieved. Peli and Lahav [5] used bicolor registration: one image is displayed on the red channel, the other on the green channel. Eaton *et al.* [14] used a similar technique, but they enhanced one image using pseudochromatic coding. In all of these applications, the operator used the vascular pattern of the retina as the main cue for registering the images. This selection is natural, because the retinal background is fairly uniform, and the vessels, occupying only a small percentage of the area, are very distinct. In fact, in many applications, manual registration is obtained by a human operator placing a cursor at the crossing or bifurcation points of corresponding vessels in the two images [4], [9], [15], but manual registration is too slow and laborious for applications that require registration of many

Manuscript received February 25, 1987; revised May 15, 1987. This work was supported in part by the National Institutes of Health under Grants R01-EY05450 and R01-EY03966.

E. Peli is with the Eye Research Institute of Retina Foundation, Boston, MA 02114, the Department of Ophthalmology, Harvard Medical School, Boston, MA 02114, and the Department of Ophthalmology, Tufts University School of Medicine, Boston, MA 01111.

R. A. Augliere is with the Eye Research Institute of Retina Foundation, Boston, MA 02114.

G. T. Timberlake is with the Eye Research Institute of Retina Foundation, Boston, MA 02114 and the Department of Ophthalmology, Harvard Medical School, Boston, MA 02114.

IEEE Log Number 8715818.

images, such as in reading analysis with the SLO, which yields hundreds of frames per minute of reading.

The registration method proposed by Barnea and Silverman [16], sequential similarity detection (SSD), is said to be up to 100 times faster than the correlation method. We developed a modification of SSD that is aimed at further improving the speed and reliability of the registration by simulating the human operator approach of detecting the vessels in the fundus images, registering them quickly, and then refining the vessels' registration.

The next section describes the standard SSD algorithm and the basic modification of feature-based SSD implemented in this study. Section III describes results of experimentation with application of the feature-based registration to retinal images obtained with different imaging modalities. Comparisons of cross-correlation, standard SSD, and feature-based SSD are included. Section IV describes further improvement in the speed and accuracy of the algorithm through incorporation of two-stage processing. Discussion of the algorithm and comparison to other results are included in Section V.

II. Feature-Based SSD

The process of matching a template in a search area is common to many registration techniques. A window $F_1(j, k)$ of size $J \times K$ is defined in one image as the template. Another window $F_2(j, k)$ of size $M \times N$ in a separate image is defined as the search area. The search area is larger than the template and is assumed to include features that resemble, but may not be identical to, those contained in the template. The template window is then shifted across the search area, and a similarity measure is calculated. The point of maximal similarity is designated as the match position, where the features within the template area are aligned with the corresponding features in the search area. The SSD developed by Barnea and Silverman [16] is an improvement of the process using the sum of the absolute values of the differences (SAVD) as the similarity measure.

$$\mathrm{SAVD}(m, n) = \sum_{j=1}^{J} \sum_{k=1}^{K} \left| F_1(j, k) - F_2(j - m, k - n) \right| \quad (1)$$

where the point with the minimal SAVD is defined as the point of registration.

The SSD, using the same similarity measure, introduced two important improvements [16]. The first is normalization, which corrects for differences in illumination between the two images. The algorithm accumulates an error of normalized absolute values of differences.

$$E(m, n) = \sum_{j} \sum_{k} \left| (F_1(j, k) - TT) - (F_2(j - m, k - m) - SS(m, n)) \right| \quad (2)$$

where TT is the average of the gray levels in the template, and SS is the local average of a section of the template size within the search area. This normalization improves the reliability of the registration.

Speed, the second important improvement, is increased by taking the summation of the cumulative error only on a subgroup of the possible $J \times K$ template points for every position (m, n). If the cumulative error $E(m, n)$ exceeds a predetermined threshold value before all the points in the window are examined, the test of that position is halted, the number of points tested up to this point (the *count*) is recorded, and the window is shifted to the next position within the search. When all positions have been examined, the position with the maximal *count* (the largest number of points that had to be examined to reach the threshold) is defined as the point of registration. Note that only a small number of points should be calculated in windows of gross miscorrespondence. This saving in the number of computations is the main benefit of the SSD algorithm.

The most computationally intensive part of the algorithm is the calculation of the normalization factor SS. Although only a small number of points must be addressed at every window position for calculating the absolute value difference, all points must be considered in the calculation of SS. Using a fast normalization procedure that incorporates a two-dimensional running average procedure [16], the number of computational operations is reduced significantly. In addition, proper programming can reduce the number of disk access operations required for this calculation [17]. We have incorporated such a fast-averaging calculation to facilitate fast searching even in large search windows.

In the original SSD algorithm [16], the points selected for calculating cumulative error in every template position are chosen randomly. This may be the best way to select those points without *a priori* knowledge of the image, but it is inefficient and time-consuming in the case of fundus photographs, which commonly contain large areas of fairly uniform brightness representing the retinal background. The contrasting retinal vessels usually occupy a much smaller portion of the total image, typically under 10 percent. The retinal vessels can be darker than the background in normal fundus photography, or brighter than the background in fluorescein angiography. The pattern of retinal vessels renders each fundus photograph unique, and thus provides the most relevant information for the registration process. We have therefore modified the SSD algorithm to select template points for calculating accumulated error from the positions of the vessels in the template window. This selection should accelerate the search process, because the gray level difference between the vessel and background is usually much larger than the difference between two noncorresponding areas in the retinal background. For these positions of the template in the search window, the error threshold is passed quickly; unlikely registration points are rejected rapidly. In addition, this procedure may be expected to sharpen the similarity measure surface, because only points of high signal-to-noise ratio are evaluated. The selection of the

comparison test points only from the vessels gives this modified SSD its enhanced feature-detection capability with regard to vessels in the fundus image.

The selection of vessel points can be done manually by the operator using a cursor driven by a graphics bit-pad, or by automatic preprocessing. In either case, the selection of those points is required only for the template and not for the larger search area. For this reason, the preprocessing overhead in selecting the template is small, especially when multiple images are being aligned with one frame. Then only one template must be preprocessed.

Automatic selection of vessel points in the template is performed using an adaptive threshold procedure. Described previously by Peli and Lahav [5], this procedure is a modification of the adaptive threshold method of Chow and Kaneko [18]. Briefly, the template image is subdivided into windows of 8 × 8 pixels, and the standard deviation of gray levels in every window is calculated. If the standard deviation exceeds a given test level, the histogram of the window is assumed to have a bimodal distribution, and the median of the gray levels is assumed to be the threshold value. The proper test level is automatically calculated to obtain more than 100 vessel points in every window of 32 × 32 pixels. The test level is found by starting at a higher integer level and lowering the level in steps of one at a time until more than 100 points are found below threshold. An initial set of incongruent points then is selected to spread them evenly across all the vessels in the template. The incongruent points are selected by subsampling the thresholded template. First, every fourth column is examined, and if a set of vessels' points are found, the central one is selected. The same procedure is repeated for every fourth row. This places the first few points on an 8 × 8 grid across the template, and usually at or near the center of the vessels (Fig. 1). This is followed by orderly selection of the residual vessel points and random selection of the rest of the template for application whenever needed.

Only vessel positions obtained by the dynamic thresholding procedure are used. The actual gray level value for each of these points used in the calculation of accumulative error is not changed.

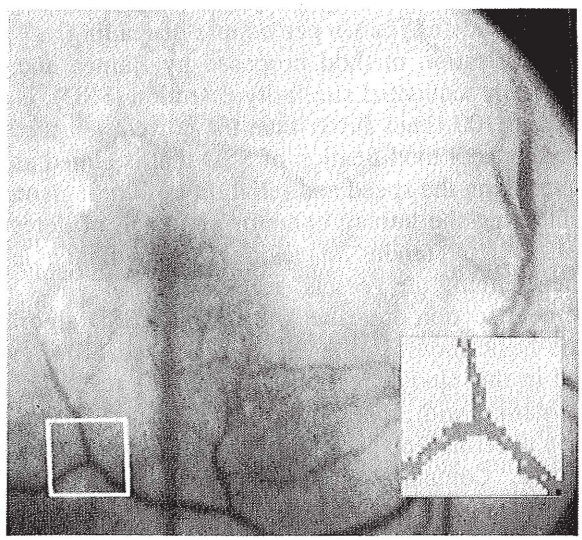

Fig. 1. An illustration of the automatic thresholding and selection of template points. The white frame marks the template area. The enlarged inset illustrates the result of automatic thresholding of this template. The bright points inside the detected vessels are the set of incongruent points that are used first in the calculation of the cumulative error.

III. Comparison of Registration Techniques

To evaluate the differences among the various registration methods, the registration programs were modified to display as a gray-scale image the similarity measure obtained with each technique. For both the standard and the feature-based SSD, the similarity measure is the *count* necessary to reach threshold in every template position within the search. Depending on the template size and the user-defined threshold of error, the *count* can cover a wide range, e.g., a template of 32 × 32 pixels might have a maximum *count* of 1024. In comparison, the cross-correlation technique produces a similarity measure that is less than 1 in absolute value. The two measures therefore must be rescaled in order to display the same intensity magnitudes for corresponding peaks of similarity. A series of different fundus images was processed using the three techniques, and the resulting similarity surface images were compared visually.

The images selected for this study include routine fundus photographs, fluorescein angiograms, and SLO video images. The pair of fundus photographs is of a patient with age-related maculopathy and drusen in the macular area (Fig. 2). One picture was taken in 1980, the other in 1983. The fluorescein angiogram pair includes two frames taken 2 s apart during the early arterial phase (Fig. 3). The SLO images were digitized from a videotape. Some of the images were separated on the tape by only a few frames, others by a number of minutes.

In all cases, a template of 32 × 32 pixels and a corresponding search area of 128 × 128 pixels were selected manually. The vessel points in the template were selected manually using the cursor and the bit-map graphics pad. The images were then processed once with the standard SSD [Figs. 2(d), 3(d)] using random selection of template points, and once with the feature-based SSD using the manually selected vessel points [Figs. 2(c), 3(c)]. The same images using the same template and similar search area then were processed using the cross-correlation process [Figs. 2(e), 3(e)].

Figs. 2 and 3 represent typical results of these experiments on two types of images. The vascular pattern of the original search area can be noted in the images representing the three similarity measures. However, the nature of those images clearly is different: the feature-based SSD results in well-delineated images of the vasculature in the search area, which is why we called it feature-based registration; however, the standard SSD results may be described as similar images obtained with multiple reflections (echoes) at different phases. The vasculature pattern also is apparent in the cross-correlation images [Figs. 2(e), 3(e)]; however, these images appear as blurred versions

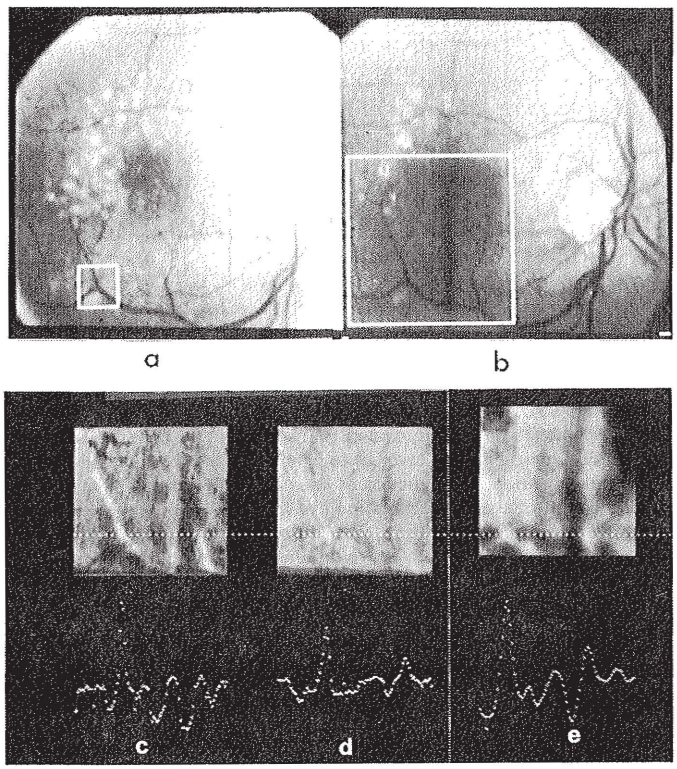

Fig. 2. Comparison of similarity measures used for registration of photographs taken three years apart. (a) Fundus photograph taken in 1980 with a 32 × 32 pixel template area marked. (b) The same eye photographed in 1983 with a 128 × 128 search area marked. (c) *Count* surface obtained with the feature-based SSD (top). Note clear appearance of the vascular pattern. The curve (bottom) represents count measure across illustrated line. Note the sharp peak at the correct registration point. (d) The *count* surface obtained with the standard (normalized) SSD. Vascular pattern is less apparent, and peak at registration is less sharp. (e) The cross-correlation surface is a blurred version of the image in (c).

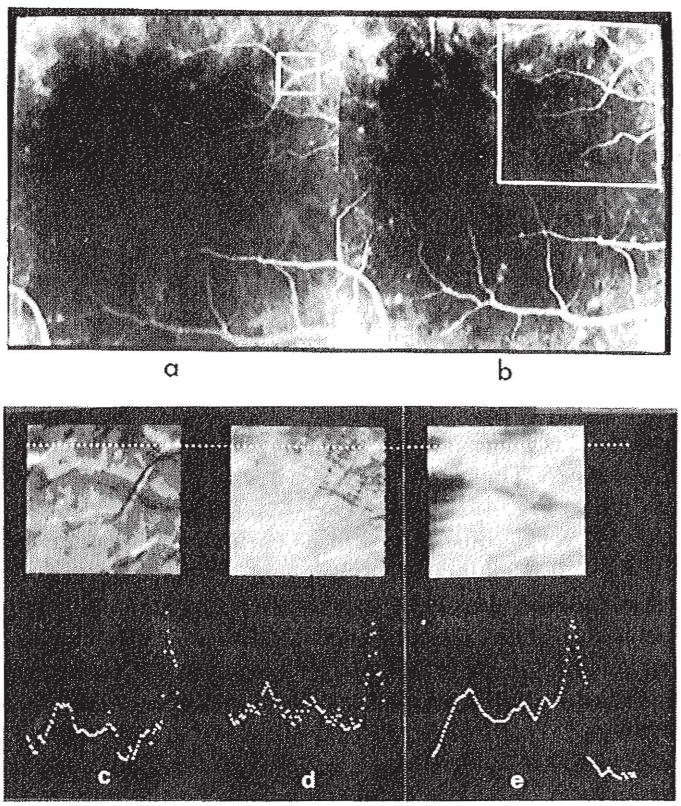

Fig. 3. Comparison of similarity measures for registration of retinal fluorescein angiograms taken 2 s apart. (a) Early photograph with template marked. (b) Later image showing increased fluorescence in the arteries (search area is marked). (c) *Count* surface obtained with the feature-based SSD. Note the vascular pattern corresponding to the search area. (d) *Count* surface obtained with the standard SSD. (e) The cross-correlation surface for the same search and template pair.

of the feature-based SSD count surfaces. The line scans taken through the maximum similarity positions illustrate the same character, i.e., smooth, less distinct peaks for the cross-correlation than for the SSD, and sharper peaks for the feature-based SSD than for the standard SSD. The line scans also illustrate that the count at the match point is higher for the feature-based SSD than for the standard SSD.

Although the similarity measure images indicate that the feature-based SSD is a sharper measure of similarity, these visual inspections do not demonstrate directly significant improvement in detection among the various techniques. The increase in speed with SSD compared with the cross-correlation process is obvious and does not require any measurement. For example, the processing of the illustrated images took more than 8 min with the cross-correlation function and less than 20 s for most SSD applications. However, for the same threshold level the modified SSD offered only a modest increase in speed of calculation.

The performances of the standard and the feature-based SSD were compared using SLO images. Templates were made from different frames from two different tapes. Another 15 to 20 frames from the same tapes were used as the search areas. Each image pair was processed twice, once with the standard SSD and once with the feature-based SSD. Each test was carried out at a number of threshold settings: 50, 100, 250, 500, 1000, 1500, and 2000. For each experiment, we recorded the completion time in seconds and the points *count* tested at the matching point. As expected, the *count* at the matching point was always higher for the feature-based than for the standard SSD. The running time for the feature-based SSD was usually shorter than the time for the standard SSD. However, the difference was only 10 to 20 percent. This difference was reduced at very high threshold levels, because more points than the 100 selected manually were required for those settings, and the additional points were selected randomly. The difference was not measurable at very low threshold settings, because our time reading was limited to a full second, with the entire calculation running 1 to 3 s.

Occasionally, the standard SSD performed even faster than the feature-based SSD, but overall the feature-based algorithm performed significantly better than the standard SSD. The difference was in the reliability of registration at low threshold levels. Our criterion for acceptable registration was a selected point within 1 pixel from the point selected by the operator by visual inspection. We noted that the feature-based SSD results were always acceptable above a certain threshold setting and not acceptable below

that setting. On the other hand, results were variable for the standard SSD. An unacceptable, erroneous registration occurred many times at a threshold higher than a lower threshold at which correct registration occurred. Normally, one would expect better detection at higher thresholds, because this offered testing of more points. Since the threshold levels were selected arbitrarily, we decided to examine further the properties of the algorithms with an orderly elevation of threshold levels.

The program was modified to run for each image pair at all integer threshold values between 1 and 300. A printed output was obtained with a threshold level, the corresponding coordinates of the selected matching point, and the point *count* tested at the selected matching point. Both fundus photographs and SLO images were tested in this manner. The results of three such testings are summarized in Fig. 4. It is noted easily that the feature-based SSD is very consistent: once correct registration is achieved, it is maintained throughout. However, for similar threshold values, the standard SSD may fluctuate between correct and grossly incorrect registration. During this testing, it also was noted that at very low threshold levels even the feature-based algorithm did not find the matching point correctly. However, the correct matching point was always within the best 10 candidates for matching point. We decided to use this valuable property in a two-stage improvement of the algorithm.

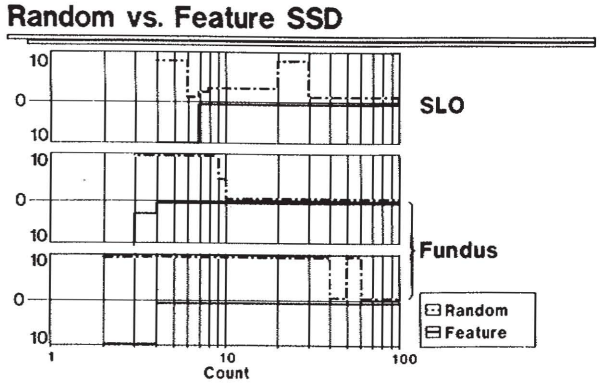

Fig. 4. Comparison of the reliability of the standard (random) and the feature-based SSD. Results from three images (one SLO image and two fundus photographs) are illustrated. In each graph the distance from the center line represents the error in registration from the visually judged correct registration, as a function of *count*. The standard SSD requires a higher count for correct registration and tends to oscillate between correct and erroneous registration. The feature-based SSD always stays in registration for *count* values above the initial *count* needed for first correct registration.

IV. Two-Stage Template Matching

Two-stage template matching was used previously to reduce calculation cost [18]–[20]. Usually the concept involves a coarse-fine two-stage approach. In the coarse stage, either a subtemplate is used for all possible positions [19], [20], or a low-resolution template is applied to a lower-resolution image of the search area [21]. The coarse stage is used to identify the "most likely" candidate positions in the search area to contain the exact match. In the second stage, the entire template is used over all possible template positions within a limited "likely" region within the search area.

The most important difficulty encountered with this coarse-fine approach is that coarse subsampling may actually skip the important points in the search area, resulting in erroneous determination in the first stage. This is especially true for a pattern with small, sharp details such as that seen in retinal vessels.

Another approach to two-stage matching involves a change of the similarity measure between the first and second stages in which a fast similarity measure is used first to determine the likely positions for a match. Then a more accurate, but slower, similarity measure could be used to select the best match among the likely ones. This approach was used previously in correlating rotated images using rotationally invariant normalized moments [22].

Both these techniques require two runs through the search area and two different sortings of the similarity measures in the two stages. The SSD algorithm enables us to implement a two-stage, two-measure process using only a single run through the search points.

Using the SSD algorithm, template positions that are likely candidates for match have a low rate of error-increase, which can be estimated simply by evaluating the *count* at various threshold levels. As each temporary threshold is passed, the *count* at that point is tested, with "high" *counts* representing likely candidates for matching. Because we found that the correct matching was included in the 10 most likely points for almost all threshold levels in the experiment described in Section III, it seemed safe to use a two threshold-two stage technique to accelerate the feature-based algorithm further.

The two-stage algorithm is applied in the following way: A low-threshold level is selected, such that the *count* at matching point and other likely points is fairly low, 1 to 2 percent of the template points. When this threshold is reached, the *count* is tested. If it exceeds 1 percent of the points, the threshold for this template position is multiplied five times, and the process continues. Thus, more template points are examined, and the error continues to accumulate. When the new higher threshold is reached, the process is completed, the low threshold is reset, and the template is shifted to the next search position. When the initial first-stage threshold is reached, if the *count* is less than 1 percent of the template points, this position is considered an unlikely match, and the template is shifted to the next position. In this process, unlikely candidate positions are determined based on evaluating less than 1 percent of the template points, which usually means use of less than 10 percent of the selected vessel points. A likely candidate for high similarity measure is being defined by elevating the threshold and continuing the SSD process. This process is summarized in the pseudocode format.

```
Next Window:        T = Threshold
                    Count = 0
                    likely = FALSE
                    Shift template to next search window.

Error Calculation:  Test the next pair of points and
                    calculate cumulative error.

                    Count = Count + 1
                      If (likely) goto Last Test

                      If error > T then
                        If count > (J × K/100) then
                          likely = TRUE
                          T = T × 5
                        Else goto Next Window
                        End If
                      Else goto Error Calculation
                      End If

Last Test:          If error > T then
                      store count
                      goto Next Window
                    Else goto Error Calculation
                    End If
```

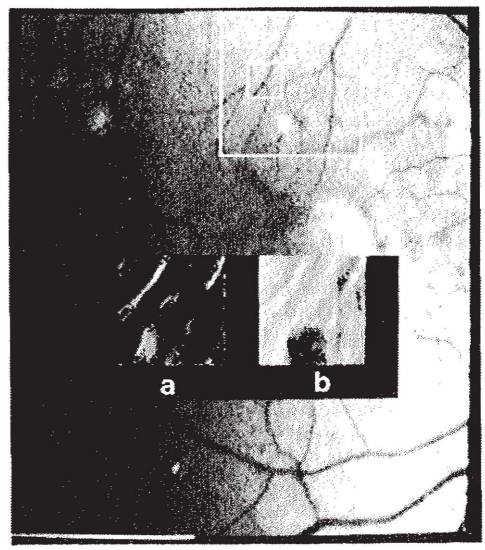

Fig. 5. Comparison of the performance of the two-stage matching with the standard and the feature-based SSD. The SLO video frame shown includes the 128 × 128 search area marked by a white square. The 32 × 32 template from another frame is shown in its original position (shifted due to eye movement). The *count* surfaces illustrated in (a) and (b) represent the results of the search with the feature-based and random algorithms, respectively. The integrated brightness of the surface is a measure of the number of calculations required for the search. The local brightness represents the time required for calculation for every template position.

This process is comparable to processing by a human operator quickly aligning the vessels and then carefully refining the position of the vessels in the two images for accurate overlap. The results of applying the two-stage algorithm in fundus images taken with the scanning laser ophthalmoscope are illustrated in Fig. 5. In the *count* surfaces illustrated, lighter points appeared when a higher *count* was achieved. The larger the darker area is in this image, the shorter the time of calculation required for registration.

V. Discussion

Image registration is a basic problem in many digital image-analysis applications. It is of special importance in medical image processing because changes over time are analyzed frequently. In a survey of registration techniques, Kashef and Sawchuk [23] concluded that there is no universal technique for solving all registration problems. The optimal solution is highly dependent on the nature of the image and system requirements. Thus the development of a special technique for registration of fundus images is appropriate. Nevertheless, the technique described here can be applied to other images having similar properties, i.e., relatively small features of distinct gray levels in an otherwise fairly uniform background. This technique can be expanded to other types of images that can be preprocessed easily to obtain such properties, e.g., Landsat images [24] processed with the gradient operator and displayed as the magnitude of the gradient image have such an appearance. Other images, e.g., chromosome images [25], are also appropriate for registration with this algorithm following gradient operation.

The SSD, or the sum of absolute value of difference (SAVD) algorithm, was compared previously to other similarity measures [24], [26]. It was inferior in both cases, because more erroneous registrations resulted than with the cross-correlation technique [24] and the stochastic sign change (SSC) measure [26]. However, the *non*-normalized version of the algorithm (1) was used in both cases rather than the normalized version (2). Although the nonnormalized version is much faster, it is prone to registration errors, especially if the images differ significantly in illumination. Even with this version, the algorithm performed well in both tests. Svedlow *et al.* [24] indicated that if computation time is examined, the nonnormalized SSD can be more advantageous than the cross-correlation algorithm. Although the normalized version of the SSD is much more intensive computationally, it is still up to 100 times faster than the correlation method if properly programmed [11], [16]. The modification of the algorithm presented here, with the ordered selection of the template points first from the vessel, adds relatively small overhead cost for calculation. The increased reliability of registration with the modified SSD enables implementation of the two-stage approach, which is necessary to increase further the speed of registration.

References

[1] P. Nagin, B. Schwartz, and G. Raynolds, "Measurement of fluorescein in angiograms of the optic disc and retina using computerized image analysis," *Ophthalmology*, vol. 92, pp. 547–552, Apr. 1985.

[2] V. R. Algazi, J. L. Keltner, and C. A. Johnson, "Computer analysis of the optic cup in glaucoma," *Invest. Ophthalmol. Vis. Sci.*, vol. 26, pp. 1759–1770, Dec. 1985.

[3] T. N. Cornsweet, S. Hersh, J. C. Humphries, R. J. Beesmer, and V. W. Cornsweet, "Quantification of the shape and color of the optic nerve head," in *Advances in Diagnostic Visual Optics (Proc. 2nd Int. Symp.)*, C. M. Breinin and I. M. Siegel, Eds. New York: Springer-Verlag, 1983.

[4] J. W. Warnicki, P. Rehkopf, J. L. Cambier, and M. R. Nelson, "Development of an imaging system for ophthalmic photography," *Biol. Photographic Ass.*, vol. 53, pp. 9–17, 1985.

[5] E. Peli and M. Lahav, "Drusen measurement from fundus photographs using computerized image analysis," *Ophthalmology*, vol. 93, pp. 1575–1580, Dec. 1986.

[6] E. Peli, T. R. Hedges, and B. Schwartz, "Computerized enhancement of retinal nerve fiber layer," *Acta Ophthalmol.*, vol. 64, pp. 113–122, Apr. 1986.

[7] P. Nagin, B. Schwartz, and K. Nanba, "The reproducibility of computerized boundary analysis for measuring optic disc pallor in the normal optic disc," *Ophthalmol.*, vol. 92, pp. 243–251, Feb. 1985.

[8] V. Miszalok, T. Seiler, and J. Wollensak, "Quantitative evaluation of long-term fundus changes," *Invest. Ophthalmol. Vis. Sci.*, vol. 25 (ARVO suppl.), p. 253, Mar. 1984.

[9] G. T. Timberlake, M. A. Mainster, E. Peli, R. A. Augliere, E. A Essock, and L. E. Arend, "Reading with a macular scotoma. I: Retinal location of scotoma and fixation area," *Invest. Ophthalmol. Vis. Sci.*, vol. 27, pp. 1137–1147, July 1986.

[10] F. Glazer, G. Raynolds, and P. Annanan, "Scene matching by hierarchical correlation," in *Proc. IEEE Comput. Soc. Conf. Computer Vision and Pattern Recognition*, 1983, pp. 432–441.

[11] W. K. Pratt, "Image registration," in *Digital Image Processing*, W. K. Pratt, Ed. New York: Wiley, 1978, pp. 562–570.

[12] A. Ling, T. Krile, S. Mitra, and Z. Shihab, "Early detection of glaucoma using digital image processing," *Invest. Ophthalmol. Vis. Sci.*, vol. 27 (ARVO suppl.), p. 160, Mar. 1986.

[13] J. A. Parker, R. V. Kenyon, and L. R. Young, "Measurement of torsion from multitemporal images of the eye using digital signal processing techniques," *IEEE Trans. Biomed. Eng.*, vol. BME-32, pp. 28–36, Jan. 1985.

[14] A. M. Eaton, D. L. Hatchell, and D. B. Chandler, "Measurements of retinal blood vessel diameter using computerized image analysis," *Invest. Ophthalmol. Vis. Sci.*, vol. 27 (ARVO suppl.), p. 307, Mar. 1986.

[15] R. W. Rowe, S. Packer, J. Rosen, and Y. Bizais, "A charge coupled device imaging system for ophthalmology," *Proc. SPIE (Application of Optical Instrumentation in Medicine XII)*, vol. 454, pp. 65–71, 1984.

[16] D. I. Barnea and H. F. Silverman, "A class of algorithms for fast digital image registration," *IEEE Trans. Comput.*, vol. C-21, pp. 179–186, 1972.

[17] T. S. Huang, G. J. Yang, and G. Y. Tang, "A fast two-dimensional median filtering algorithm," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-27, pp. 12–18, 1979.

[18] C. K. Chow and T. Kaneko, "Automatic boundary detection of the left ventricle from cineangiograms," *Comput. Biomed. Res.*, vol. 5, pp. 388–410, Aug. 1972.

[19] G. J. Vanderburg and A. Rosenfeld, "Two-stage template matching," *IEEE Trans. Comput.*, vol. C-26, pp. 384–393, 1977.

[20] A. Goshtasby, S. H. Gage, and J. F. Bartholic, "A two-stage cross correlation approach to template matching," *IEEE Trans. Pattern Anal. Machine Intell*, vol. PAMI-6, pp. 374–378, 1984.

[21] A. Rosenfeld and G. J. Vanderburg, "Coarse-fine template matching," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-2, pp. 104–107, 1977.

[22] G. Ardeshir, "Template matching in rotated images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, no. 3, pp. 338–344, 1985.

[23] B. G. Kashef and A. A. Sawchuk, "Survey of new techniques for image registration and mapping," *Proc. SPIE* (Application of Digital Image Processing VI), vol. 432, pp. 222–239, 1983.

[24] M. Svedlow, C. D. McGlem, and P. E. Anuta, "Image registration: Similarity measure and processing method comparisons," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-14, no. 1, pp. 141–149, 1978.

[25] L. Kitchen and A. Rosenfeld, "Grey level corner detection," Dep. Comput. Sci., Univ. Maryland, College Park, Tech. Rep. TR-887, 1980.

[26] A. Venot and E. Leclerc, "Automated correction of patient motion and grey values prior to subtraction in digitized angiography," *IEEE Trans. Med. Imaging*, vol. MI-3, no. 4, pp. 179–186, 1984.