(19) DANMARK

(10) DK/EP 3449464 T3

(12) Oversættelse af europæisk patentskrift

Patent- og Varemærkestyrelsen

(51) Int.Cl.: *G 06 T 7/73 (2017.01)*

(45) Oversættelsen bekendtgjort den: **2020-05-11**

(80) Dato for Den Europæiske Patentmyndigheds bekendtgørelse om meddelelse af patentet: **2020-04-08**

(86) Europæisk ansøgning nr.: **17720737.0**

(86) Europæisk indleveringsdag: **2017-04-24**

(87) Den europæiske ansøgnings publiceringsdag: **2019-03-06**

(86) International ansøgning nr.: **EP2017059647**

(87) Internationalt publikationsnr.: **WO2017186634**

(30) Prioritet: **2016-04-26 US 201615139179**

(84) Designerede stater: **AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**

(73) Patenthaver: **Optos Plc, Queensferry House , Carnegie Campus , Enterprise Way, Dunfermline, Scotland, KY11 8GR, Storbritannien**

(72) Opfinder: **FLEMING, Alan Duncan, 14 Oxgangs Hill, Edinburgh, Scotland EH13 9JR, Storbritannien**

(74) Fuldmægtig i Danmark: **Plougmann Vingtoft A/S, Strandvejen 70, 2900 Hellerup, Danmark**

(54) Benævnelse: **BEHANDLING AF NETHINDEBILLEDE**

(56) Fremdragne publikationer:
ALAN D FLEMING ET AL: "Automatic detection of retinal anatomy to assist diabetic retinopathy screening", PHYSICS IN MEDICINE AND BIOLOGY, INSTITUTE OF PHYSICS PUBLISHING, BRISTOL GB, vol. 52, no. 2, 21 January 2007 (2007-01-21), pages 331-345, XP020113152, ISSN: 0031-9155, DOI: 10.1088/0031-9155/52/2/002
G Linda Shapiro ET AL: "Computer VIsion - chapter 10 Image Segmentation", , March 2000 (2000-03), XP055199831, Retrieved from the Internet: URL:https://courses.cs.washington.edu/courses/cse576/book/ch10.pdf [retrieved on 2015-07-02]

DESCRIPTION

[Technical Field]

[0001] The present invention generally relates to the field of medical image processing and, more particularly, to the processing of a retinal image to estimate the location of the fovea of the retina within the image.

[Background]

[0002] Many eye diseases can be detected and diagnosed by imaging the retina. In a busy clinic it is not always possible to record even basic information such as whether the image is of the left or right eye. Automated analysis and manipulation of retinal images is a growing field. Automated analysis by computerised algorithms can provide assistance to clinicians in the form of disease detection or improved presentation modes. Many automated methods require positional information of the basic anatomical landmarks to aid their operation. This positional information may be input into software which may, for example: (i) project a retinal image into a uniform format which allows accurate dimensional measurements of the retina to be made; (ii) analyse multiple retinal field-of-view image sets, for example to identify fields-of-view and confirm that at all required fields-of-view are present; (iii) register multiple retinal images so that the location of disease visible in one modality can be identified in another; (iv) montage multiple retinal images to present clinicians with a wider view of the retina in one image; (v) present the user with a fly-through presentation of the retina; and/or (vi) perform automated disease determination; this can be so that the software operates only on the relevant area (for example, the optic disc in glaucoma, the macula in age-related macular degeneration) or so that the software ignores irrelevant areas (for example, the optic disc in diabetic retinopathy).

[0003] The article titled "Automatic detection of retinal anatomy to assist diabetic retinopathy screening" by A. D. Fleming et al., published in Phys. Med. Biol. 52 (2007) 331-345 on 21 January 2007, discloses methods for the robust location of the optic disc and fovea. The elliptical form of the major retinal blood vessels is used to obtain approximate locations, which are refined based on the circular edge of the optic disc and the local darkening at the fovea. The methods have been tested on 1056 sequential images from a retinal screening programme. Positional accuracy was better than 0.5 of a disc diameter in 98.4% of cases for optic disc location, and in 96.5% of cases for fovea location. The methods are sufficiently accurate to form an important and effective component of an automated image grading system for diabetic retinopathy screening.

[0004] Reference is also directed to "Computer Vision - Chapter 10 Image Segmentation" by G. L. Shapiro et al., March 2000.

[0005] Although methods have been developed for determining the locations of anatomical features of the retina such as the optic disc and fovea in fundus reflectance photographs (mainly red and green light), the technical problem of automatically and reliably estimating the location of the fovea in a retinal images of other modalities, and for ultra-wide field views, presents a difficult challenge, owing to artefacts and pathology, variations in eye position and variations between image modalities.

[Summary]

[0006] In view of the shortcomings of conventional retinal image-processing methods and devices summarised above, the present inventor has devised a non-transitory computer-readable storage medium storing computer program instructions which, when executed by a processor, cause the processor to process image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea of the retina in the image by: receiving the image data; transforming the received image data such that the vascular structure in the image defined by the transformed image data is more circular than the vascular structure in the image defined by the received image data; calculating, for each of a plurality of pixels of the transformed image data, a respective local orientation vector indicative of the orientation of any blood vessel present in the image at or adjacent to the pixel; calculating a normalised local orientation vector for each of the plurality of pixels by normalising the local orientation vector calculated for each of the pixels so that the magnitude of the normalised local orientation vector at each of the pixels takes a common predetermined value; operating on an array of accumulators, wherein each accumulator in the array is associated with a respective pixel of the transformed image, by (i) for each of the plurality of pixels, adding a respective value to an accumulator associated with a pixel of the transformed image data that is offset from the pixel of the plurality of pixels by the predetermined value in a predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, and (ii) smoothing the values in the accumulators; and estimating the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator having accumulated an accumulated value which is within a predetermined range of an extremum of the accumulated values in the accumulators.

[0007] The inventor has further devised an image-processing apparatus comprising a processor and a memory, the memory storing instructions executable by the processor whereby the processor is operative to process image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea of the retina in the image by: receiving the image data; transforming the received image data such that the vascular structure in the image defined by the transformed image data is more circular than the vascular structure in the image defined by the received image data; calculating, for each of a plurality of pixels of the transformed image data, a respective local orientation vector indicative of the orientation of any blood vessel present in the image at or adjacent to the pixel; calculating a normalised local orientation vector for each of the plurality of pixels by normalising

the local orientation vector calculated for each of the pixels so that the magnitude of the normalised local orientation vector at each of the pixels takes a common predetermined value; operating on an array of accumulators, wherein each accumulator in the array is associated with a respective pixel of the transformed image data, by (i) for each of the plurality of pixels, adding a respective value to an accumulator associated with a pixel of the transformed image data that is offset from the pixel of the plurality of pixels by the predetermined value in a predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, and (ii) smoothing the values in the accumulators; and estimating the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator having accumulated an accumulated value which is within a predetermined range of an extremum of the accumulated values in the accumulators.

[0008] In an embodiment of the image-processing apparatus, the instructions, when executed by the processor, may cause the processor to perform, prior to transforming the received image data, further processes of: (i) processing the received image data to generate lower spatial resolution image data defining a lower spatial resolution image of the vascular structure of the temporal vascular arcades of the retina; and (ii) generating filtered image data by filtering the lower spatial resolution image data so as to enhance linear structures in the image defined by the lower spatial resolution image data. The processor may transform the filtered image data such that the vascular structure in the image defined by the transformed image data is more circular than the vascular structure in the image defined by the received image data.

[0009] In this embodiment, the instructions, when executed by the processor, may cause the processor to perform, at least once, prior to transforming the received image data, further processes of: (iii) processing the filtered image data to generate lower spatial resolution filtered image data defining a low spatial resolution image of the vascular structure of the temporal vascular arcades of the retina; (iv) filtering the lower spatial resolution filtered image data so as to enhance linear structures in the image defined by the lower spatial resolution filtered image data, wherein the filtered image data processed in any repeat of the process (iii) is the filtered image data generated in the preceding performance of process (iv), and wherein the filtered image data transformed by the processor is the filtered image data generated in the final performance of process (iv). The instructions when executed by the processor, may cause the processor to downsize the received image data by a factor of 1.4 to 4 in the performance of process (i), and to downsize the filtered image data by a factor of 1.4 to 4 in each performance of process (iii).

[0010] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to operate on a plurality, N, of arrays of accumulators, where N is an integer equal to at least 2, and wherein: each of the arrays of accumulators is associated with a respective one of N quantised directions for the local orientation vector; and each of the accumulators in each of the N arrays is associated with a respective one of the pixels in the

transformed image data. The instructions, when executed by the processor, may cause the processor to operate on the plurality of arrays of accumulators by, for each of the plurality of pixels: quantising the direction of the calculated local orientation vector at the pixel into one of the N quantised directions; and adding a respective value to an accumulator of an accumulator array that is associated with the same quantised direction as the quantised direction of the pixel value gradient vector at the pixel, the accumulator being associated with a pixel of the transformed image data that is offset from the pixel of the plurality of pixels by the predetermined value in the predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels; and by smoothing the values in the accumulators of the plurality of arrays.

[0011] The instructions, when executed by the processor, may further cause the processor to: further calculate, for each set of N accumulators that are provided in different respective accumulator arrays but are associated with the same pixel of the transformed image data, a product of the respective accumulated values that have been accumulated in the N accumulators; and estimate the location of the fovea in the image of the retina using the location of a pixel of the transformed image data associated with one of the sets of N accumulators, for which set of accumulators the calculated product of accumulated values is within a predetermined range of an extremum of the calculated products of accumulated values. Preferably, $2 \leq N \leq 32$, and more preferably $4 \leq N \leq 16$.

[0012] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to calculate a second normalised local orientation vector for each of the plurality of pixels by normalising the local orientation vector calculated for each of the plurality of pixels so that the magnitude of the second normalised local orientation vector at each of the pixels takes a second common predetermined value, and to further operate on the array of accumulators by: for each pixel of the plurality of pixels, adding a respective value to an accumulator associated with a pixel in the transformed image data that is offset from the pixel of the plurality of pixels by the second predetermined value in the predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels.

[0013] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to process the image data to estimate the location of the fovea of the retina in the image by performing a further process of: selecting pixels from the plurality of pixels of the transformed image data such that the magnitude of the calculated local orientation vector at each of the selected pixels exceeds a threshold, and by: calculating the normalised local orientation vector for each of the selected pixels by normalising the local orientation vector calculated for each of the selected pixels so that the magnitude of the normalised local orientation vector at each of the selected pixels takes a common predetermined value; and operating on the array of accumulators by: for each of the selected pixels, adding a respective value to an accumulator associated with a pixel of the transformed image data that is offset from the selected pixel in the image by the predetermined value in the predetermined direction

relative to the direction of the local orientation vector at the selected pixel. In this case, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to select the pixels from the plurality of pixels such that the magnitude of the calculated local orientation vector at each of the selected pixels exceeds the 50^{th} percentile, and more preferably the 80^{th} percentile, of local orientation vector magnitudes of the pixels in the image.

[0014] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to operate on each array of accumulators by: for each pixel of the plurality of pixels, adding, as the respective value, a respective weighting to the accumulator associated with a pixel in the transformed image data that is offset from the pixel of the plurality of pixels by the predetermined value in the predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, the weighting being indicative of the magnitude of the local orientation vector at the pixel of the plurality of pixels.

[0015] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to calculate the normalised local orientation vector for each pixel of the plurality of pixels so that that the magnitude of the normalised local orientation vector at each pixel of the plurality of pixels is between 0.5 DFD and 1.3 DFD, where DFD is the distance between the fovea and the center of an optic disc in the image defined by the received image data.

[0016] In the image-processing apparatus or any of the embodiments set out above, the instructions, when executed by the processor, may, in accordance with a further embodiment, cause the processor to smooth the values in the accumulators using a kernel having a standard deviation of 0.02 DFD to 0.5 DFD, and more preferably 0.075 DFD to 0.25 DFD, where DFD is the distance between the fovea and the center of an optic disc in the image defined by the received image data.

[Brief Description of the Drawings]

[0017] Embodiments of the invention will now be explained in detail, by way of example only, with reference to the accompanying figures, in which:

Fig. 1 is a schematic illustration of an image-processing apparatus according to embodiments of the present invention;

Fig. 2 is a first part of a flow diagram illustrating a process by which the image-processing apparatus of the first embodiment processes image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea in the image;

Figs. 3A to 3C show results of filtering operations performed in step S20 in Fig. 2;

Fig. 4 is a continuation of the flow diagram in Fig. 2;

Fig. 5 is a flow diagram illustrating process S70 in Fig. 4;

Fig. 6 is a schematic illustrating the relationship between pixels p^{s}_{ij} and p^{o}_{ij} of a retinal image and an accumulator a_{ij} of an accumulator array described herein;

Fig. 7 is an alternative continuation of the flow diagram in Fig. 2, which illustrates a process by which the image-processing apparatus according to a second embodiment processes image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea in the image; and

Fig. 8 is another alternative continuation of the flow diagram in Fig. 2, which illustrates a process by which the image-processing apparatus according to a variant of the first embodiment processes image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea in the image.

[Detailed Description of Embodiments]

[Embodiment 1]

[0018] Figure 1 illustrates an image-processing apparatus according to a first embodiment, which is configured to process image data defining an image of a vascular structure of the temporal vascular arcades of a retina to estimate a location of the fovea of the retina in the image. The two temporal vascular arcades (namely, the superior and inferior temporal vascular arcades, each comprising a vein and an artery) found in the retina of the eye define a vascular structure which follows a generally elliptical path as the blood vessels branch out from the optic disc, with the macula of the eye being surrounded by this elliptical path, and the center of the macula (namely, the fovea) being located approximately at the centre of the ellipse. The image-processing apparatus of the present embodiment exploits the geometry of this vascular structure in retinal images, using it to estimate the location of the fovea, by determining the location of a pixel which is indicative of the location of the fovea.

[0019] The image-processing apparatus of the present embodiment forms part of a programmable signal processing apparatus, as illustrated schematically in Fig. 1. The signal processing apparatus 100 shown in Fig. 1 comprises an input/output (I/O) section 110 for receiving image data defining an image of the retina, and for outputting the determined location of the pixel which is indicative of the location of the fovea, or a result of further

processing operations based thereon, for example: (i) a more refined estimate of where the fovea is likely to be (which may correspond to the location of a pixel that is separated from the aforementioned pixel by a predetermined displacement vector); (ii) a determination of whether the image is of the right eye or the left eye; and/or (iii) an estimate of where the macula lies in the image. The determination of whether the image is of the right eye or the left eye may be determined by comparing the determined location of the pixel which is indicative of the location of the fovea with, for example, the center of the image, with a determined location of the optic disc, or with the location of a prominent darkening in the image, using techniques known to those skilled in the art.

[0020] The signal processing apparatus 100 further comprises a processor 120, a working memory 130 and an instruction store 140 storing computer-readable instructions which, when executed by the processor 120, cause the processor 120 to perform the processing operations hereinafter described to determine the location of the fovea in the retinal image. The I/O section 110, the working memory 130 and the instruction store 140 are communicably connected to the processor 120, as shown in Fig. 1. The instruction store 140 may comprise a ROM which is pre-loaded with the computer-readable instructions. Alternatively, the instruction store 140 may comprise a RAM or similar type of memory, and the computer-readable instructions can be input thereto from a computer program product, such as a computer-readable storage medium 150 such as a CD-ROM, etc. or a computer-readable signal 160 carrying the computer-readable instructions. The signal processing apparatus 100 may, for example, form part of a personal computer (PC) or a network server, and may comprise other components that are well-known (e.g. display, input device(s), etc.). In the present embodiment, the combination 170 of the hardware components shown in Fig. 1, comprising the processor 120, the working memory 130 and the instruction store 140, is configured to implement the functionality of the image-processing apparatus of the present embodiment, which will now be described in detail with reference to Figs. 2 to 6.

[0021] Figure 2 is a flow chart illustrating a first part of a process by which the image-processing apparatus 170 processes image data defining an image of a vascular structure of temporal vascular arcades of a retina to estimate a location of the fovea of the retina, which may be useful in a variety of different applications.

[0022] Firstly, in step S10, the processor 120 receives image data defining an image of a vascular structure of temporal vascular arcades of a retina via the I/O section 110. The acquired image data may define an retinal image which may be one of a number of different types that are known to those skilled in the art, for example a red and green (RG) reflectance image, an autofluorescence (AF) image, a fluorescein angiogram (FA) image, a reflectance scanning laser ophthalmoscope (SLO) image of the retina, or a colour fundus photograph of the retina, and may provide an ultra-wide-field view of the retina.

[0023] The processor 120 may then proceed to generate filtered image data in an optional step S20, wherein the processor 120 may pre-process the received retinal image data to enhance the vasculature in the image, using techniques known to those skilled in the art, for

example as described in chapter 8 of the book "Automated Image Detection of Retinal Pathology" by H. Jelinek and M. J. Cree (CRC Press, January 1, 2009). For example, the complex-valued Gabor kernel

$$g(x,y;\lambda,\theta,\psi,\sigma,\gamma)=\exp\left(-\frac{x'^2+\gamma^2y'^2}{2\sigma^2}\right)\exp\left(i\left(2\pi\frac{x'}{\lambda}+\psi\right)\right)$$

may be used to generate convolved images, with the value of θ taking each value $2a\pi/N$ for a = 0... N - 1, where N is the number of orientations used in the evaluation which, may, for example, be between 8 and 12 (although reducing N to 6 in the pre-processing of RG images was not found to be detrimental). Values used for the other parameters in the Gabor kernel g may, for example, be:

$\sigma = 2.75\sqrt{s},\ \gamma = 0.37, \lambda = 7s, \psi = 0.$

[0024] The parameter s sets the scale. The value of s = 2 was found to be well-suited for the pre-processing of Optomap™ images. However, the ideal value of s depends on the magnification of the image (or the real dimension in microns per pixel) and this is variable between imaging modalities and image formats. Also, a range of values of s may be used to account for vessels with a range of diameters. In this case, the resulting vessel enhancements at each value of s may be combined, for example, by pixel-by-pixel addition.

[0025] In this example, an array of complex-valued images is generated by convolution of the image by each Gabor kernel. Each element of this array is generated with one of the kernels:

$$\mathrm{U(a)} = -\mathrm{I} \circ \mathrm{g}\left(\lambda, \frac{\mathrm{a}\pi}{N}, \psi, \sigma, \gamma\right)$$

[0026] At each pixel, the orientation index is found which gave the maximum absolute response:

$$\mathrm{A(x,y)} = \arg\max_{\mathrm{a}=0\ldots\mathrm{N}-1} |U(x,y;a)|$$

[0027] This is the locally dominant (discretised) angle for pixels on a vessel, i.e. the orientation of the vessel.

[0028] The complex vessel-enhanced image, V, in which blood vessels appear brighter than the background, has a value at each pixel which is selected from one of the images in array U:

$$\mathrm{V(x,y)} = \mathrm{U(x,y;A(x,y))},\ \forall(x,y)$$

[0029] Another form of linear filtering, non-linear filtering and/or normalisation may alternatively be applied at the pre-processing stage. For example, a low-pass filter may be used to reduce noise. Additionally or alternatively, a high-pass filter may be used to remove gradual image intensity changes such as lighting variations. A high-pass filter may also aid more consistent image normalisation. Normalisation is any method serving to reduce the variation between

images of some parameter such as global mean brightness, global mean contrast or image histogram, or to reduce the variation across a single image of parameters such as local image brightness, local contrast or local histogram.

[0030] The processor 120 may, as in the present embodiment, generate the filtered image data in step S20 firstly by processing the received image data to generate lower spatial resolution image data defining a lower spatial resolution image of the vascular structure of the temporal vascular arcades of the retina. This may be done in any suitable way using techniques known to those skilled in the art, for example by appropriately sampling the received image data and (if necessary) applying any desired form of interpolation (e.g. bilinear interpolation), spatial anti-aliasing, etc. The processor 120 may reduce the spatial resolution of the received image data by a factor of 1.4 to 4, and more preferably 2 (as in the present embodiment). The processor 120 may then generate the filtered image data by filtering the generated lower spatial resolution image data so as to enhance linear structures in the image defined by the lower spatial resolution image data using, for example, the techniques described above.

[0031] The processor 120 may then perform further processes of: (i) processing the filtered image data to generate lower spatial resolution filtered image data defining a low spatial resolution image of the vascular structure of the temporal vascular arcades of the retina (the spatial resolution again preferably being reduced by a factor of 1.4 to 4, more preferably 2); and (ii) filtering the lower spatial resolution filtered image data generated in process (i) so as to enhance linear structures in the image defined by the lower spatial resolution filtered image data, wherein the filtered image data processed in any repeat of the process (i) is the filtered image data generated in the preceding performance of process (ii). In this way, processes (i) and (ii) may be performed iteratively, two or more times (and preferably twice), with the output of process (ii) in one iteration serving as the input to process (i) in the next iteration. Examples of the initial image enhancement and two successive repeats of processes (i) and (ii) are shown in Figs. 3A, 3B and 3C, respectively, in which the resulting smoothed images have been sized to the same dimensions.

[0032] The process then proceeds to step S30, wherein the processor 120 transforms the filtered image data generated at the end of the above-described processes in step S20 (or the received image data in embodiments where step S20 is omitted) such that the vascular structure in the image defined by the transformed image data is more circular than the vascular structure in the image defined by the filtered image data (or received image data, as the case may be). This transformation may be done by lowering the resolution of the filtered image data in either the horizontal or the vertical direction of the image by an amount required to make the vascular structure more circular in the image, using techniques known to those skilled in the art. Making the vascular structure in the image more circular allows the simplified processing operations described herein below to be used, rather than more complex and computationally expensive calculations that would be required to estimate the location of the fovea in an image wherein the vascular structure is elliptical.

[0033] In step S40 in Fig. 2, the processor 120 calculates, for each of a plurality of pixels p_{ij} of the transformed image (where i and j represent the column and row of the pixel array, respectively), a respective local orientation vector indicative of the orientation of any blood vessel present in the image at or adjacent to the pixel (in other words, any blood vessel that may be present in the image at, or in the vicinity of (e.g. within a predetermined number of one or more pixels from), the pixel). Such a local orientation vector may be found in a number of different ways, using techniques known to those skilled in the art. Local orientation and edge strength can be evaluated using first-order or second-order derivatives of image intensity, for example. First-order derivatives can be calculated by convolution with a Sobel or other kernel (as discussed above). Local orientation is then the arctangent of the ratio of the y-component and the x-component of the gradient. Second-order derivatives can be calculated by repeated convolution with a Sobel or other kernel. The local orientation is then the arctangent of the y-component and the x-component of the eigenvector with the largest eigenvalue of the Hessian matrix:

$$\begin{bmatrix} \frac{d^2z}{dx^2} & \frac{d^2z}{dydx} \\ \frac{d^2z}{dxdy} & \frac{d^2z}{dy^2} \end{bmatrix},$$

where z is image intensity.

[0034] For example, the local orientation vector may be calculated by finding the eigenvector of the Hessian matrix at each pixel that has the largest eigenvalue. Alternatively, the local orientation vector may, as in the present embodiment, take the form of the pixel value gradient vector, which tends to have a large magnitude in the boundary region of a blood vessel in the retinal image, where the pixel values change relatively rapidly one pixel to the next in the direction perpendicular to the local orientation of the blood vessel segment (i.e. along or opposite to the flow direction of the blood vessel; in other words, the direction along which the blood vessel extends in the image), and is generally small away from the blood vessels. The pixel value gradient vector at a pixel in a region of the retinal image showing a part of a blood vessel (particularly an edge portion of the blood vessel) is thus indicative of the orientation of the blood vessel.

[0035] Thus in the present embodiment, the processor 120 calculates in step S40, for each of the plurality of pixels p_{ij} of the transformed image data, a respective pixel value gradient vector $\mathbf{g}(i,j)$ (in other words, the gradient of image intensity) at the pixel. The pixels may be considered to hold respective values of a discrete function f(i,j) whose value varies with position (i,j) on the array of pixels, with the pixel value gradient vector **g** at any pixel pointing in the direction of greatest increase of the function at that pixel, the magnitude of the pixel value gradient vector, $|\mathbf{g}|$, being the slope of the function in that direction. As noted above, in an image of a retina, the magnitude of the pixel value gradient vector will generally be highest at the edges of a blood vessel, where there is a rapid change in pixels values in the direction perpendicular to the direction along which the blood vessel extends.

[0036] Any well-known technique for calculating the pixel value gradient vector **g** may be used, for example as described in "Digital Image Processing" by R. C. Gonzalez and R. E. Woods

(Pearson, August 31, 2007) or "Digital Image Processing and Computer Vision" by R. J. Schalkoff (John Wiley & Sons, September 2, 1992). Typical methods use Sobel, Prewitt or Roberts operators, which can be used to determine the gradient x- and y-components, the gradient magnitude and the gradient orientation. The plurality of pixels p_{ij} at which the pixel value gradient vector $\mathbf{g}(i,j)$ is calculated in step S40 may encompass all of the pixels of the transformed image data, or only some of those pixels. The processor 120 may, for example, employ a mask defining one or more regions of the transformed image (e.g. a peripheral region of the image), in which region(s) no calculation of the pixel value gradient vector **g** is to be performed.

[0037] The process may, as in the present embodiment, then proceed to an optional step S50, wherein the processor 120 selects pixels p^s_{ij} from the plurality of pixels p_{ij} such that the magnitude of the local orientation vector (in this embodiment, the calculated pixel value gradient vector $\mathbf{g}(i,j)$) at each of the selected pixels p^s_{ij} exceeds a threshold. This optional step disqualifies pixels at which the local orientation vector (i.e. $|\mathbf{g}|$) is small (i.e. pixels unlikely to be located within a blood vessel or at a blood vessel edge) from the further processing operations described below, thereby saving computational resources and, moreover, allowing the location of the fovea to be estimated with higher accuracy. The processor 120 may, for example, select the pixels p^s_{ij} from the plurality of pixels p_{ij} such that the magnitude of the calculated pixel value gradient vector $\mathbf{g}(i,j)$ at each of the selected pixels p^s_{ij} exceeds a predetermined percentile of pixel value gradient vector magnitudes of the pixels p_{ij}. The predetermined percentile may be the 50^{th} percentile, or it may, as in the present embodiment, more preferably be the 80^{th} percentile.

[0038] Proceeding to step S60 in Fig. 4, the processor 120 calculates a normalised local orientation vector for each of the selected pixels p^s_{ij} by normalising the local orientation vector calculated for each of the selected pixels p^s_{ij} so that the magnitude of the normalised local orientation vector at each of the selected pixels p^s_{ij} takes a predetermined value that is the same for all of the selected pixels p^s_{ij}. Where step S50 is omitted, the processor 120 calculates in step S60 a normalised local orientation vector for each of the plurality of pixels p_{ij} by normalising the local orientation vector calculated for each of the pixels p_{ij} so that the magnitude of the normalised local orientation vector at each pixel takes a predetermined value that is the same for all of the pixels.

[0039] As the local orientation vector is, by way of example, the pixel value gradient vector in the present embodiment, the processor calculates in step S60 a normalised pixel value gradient vector $\mathbf{g}^n(i,j)$ for each of the selected pixels p^s_{ij} by normalising the pixel value gradient vector $\mathbf{g}(i,j)$ calculated for each of the selected pixels p^s_{ij} so that the magnitude of the normalised pixel value gradient vector, $|\mathbf{g}^n(i,j)|$, at each of the selected pixels p^s_{ij} takes a

predetermined value that is the same for all of the selected pixels p^s_{ij}. Thus, $|\mathbf{g}^n(i,j)|$ is the same for all of the selected pixels p^s_{ij}, while the orientation of $\mathbf{g}^n$ will depend on the local pixel value gradient at each selected pixel p^s_{ij} and will therefore generally differ between the selected pixels. The processor 120 preferably calculates the normalised pixel value gradient vector $\mathbf{g}_n$ (or other form of local orientation vector, as noted above) for each of the selected pixels p^s_{ij} so that that the magnitude of said vector (i.e. R) at each of the selected pixels p^s_{ij} is between 0.5 DFD and 1.3 DFD, where DFD is the distance between the fovea and the center of the optic disc in the received image of the retina (expressed in terms of a number of pixels, for example). The value of DFD may readily be determined (e.g. by inspecting an image of the retina), and subsequently used in the automatic processing of the same image or other images that have been obtained under similar conditions, using the techniques described herein. In other words, a suitable value of the unit DFD to be used for automatically processing a batch of retinal images may be determined by manually examining a set of training images which are representative of the images in the batch to obtain a suitable measure of DFD, which could then be used in the processing of the batch of images.

[0040] In process S70, the processor 120 performs operations described hereinafter using an array of accumulators. Each accumulator, a_{ij}, of the array is configured to receive real values (which may or may not be integers, and may be positive or negative) and calculate an accumulated value that is indicative of an accumulation of the received values. For example, each accumulator a_{ij} may, as in the present embodiment, sum positive integer values (also referred to hereinafter as "votes") which it receives such that the accumulated value is a sum of the received values. Each accumulator a_{ij} may add each received value to a sum of any previously received values calculated thereby (in other words, the accumulator a_{ij} may update its accumulated value on receiving each new value), or it may temporarily store the received values before calculating their sum at the end of the operations in step S70. However, in other embodiments, each accumulator a_{ij} may alternatively be initialised to store an initial value which is a sufficiently large positive number, and add subsequently received negative integer values to the stored value (or deduct received positive values from the stored value), thereby effectively decrementing the stored value with each value it receives. In this case, the accumulated value at the end of the operations in step S70 is also indicative of an accumulation of the received values.

[0041] In the present embodiment, the array of accumulators is implemented by the processor 120 executing the aforementioned instructions in the instruction store 140 to appropriately address and manage (including writing to, reading from and otherwise processing information stored in) storage elements in the working memory 130. Each accumulator a_{ij} is associated with a respective pixel p_{ij} of the received image data, for example by the processor 120 in the working memory 130 an association, link or pointer relating each accumulator a_{ij} to a corresponding pixel p_{ij} in the image data.

[0042] In step S70, the processor 120 operates on the array of accumulators by adding, for each pixel of the plurality of pixels, a respective value to an accumulator a_{ij} associated with a pixel p^{o}_{ij} of the transformed image data that is offset (i.e. separated by a displacement vector) from the pixel of the plurality of pixels by the predetermined value in a predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, and by smoothing the values in the accumulators a_{ij}. The process S70 will now be described in more detail with reference to Figs. 5 and 6.

[0043] Referring firstly to Fig. 5, in step S71, the processor initialises the array of accumulators so that the accumulated value stored in each accumulator a_{ij} is common constant value; in this embodiment, zero. In the aforementioned alternative embodiments, where the accumulators subtract received positive values from (or, equivalently, add received negative values to) an initial (starting) value, each accumulator a_{ij} is set to store the initial value in step S71.

[0044] In step S72, an index I, which is used to reference each of the selected pixels p^{s}_{ij} in turn, is set to an initial value of 1.

[0045] Then, in step S73, the processor 120 processes a first of the selected pixels p^{s}_{ij} (as I = 1 at this stage) by adding a value (which, by way of example is 1 in the present embodiment) to an accumulator a_{ij} of the accumulator array which is associated with a pixel p^{o}_{ij} of the received image data that is offset (i.e. spaced apart) from the selected pixel p^{s}_{ij} in the image data by the predetermined value R in a predetermined direction relative to the direction of the local orientation vector at the selected pixel. As the local orientation vector is the pixel value gradient vector $\mathbf{g}(i,j)$ in this embodiment, the predetermined direction is parallel to direction of the pixel value gradient vector $\mathbf{g}(i,j)$ at the selected pixel p^{s}_{ij}. In other embodiments, the predetermined direction may be perpendicular to the direction of the local orientation vector, or be at another predefined angle relative to the direction of the local orientation vector, depending on how the local orientation vector is calculated. The offset considered in step S73 will generally be in a direction perpendicular to the direction in which the blood vessel illustrated by the selected pixel extends at the selected pixel, as illustrated in Fig. 6.

[0046] More particularly, Fig. 6 illustrates the relationship between pixels p^{s}_{ij} and p^{o}_{ij} of the retinal image data and an accumulator a_{ij} of the accumulator array A. In Fig. 6, the transformed image data is illustrated in the upper part of the figure, and defines an image of a retina having an optic disc OD and temporal vascular arcades TVA-1 and TVA-2, which converge on the optic disc OD via respective paths that are substantially arcs of a circle centered on the fovea (not shown in Fig. 6) in the transformed image of the retina shown in Fig. 6. The lower part of Fig. 6 illustrates a part of the accumulator array A comprising the aforementioned accumulator a_{ij}. As illustrated in Fig. 6, pixel p^{o}_{ij} is taken to be a pixel which lies on a line that is parallel to the direction of $\mathbf{g}(i,j)$ at p^{s}_{ij}, and whose distance from p^{s}_{ij} along

that line is closest to R (in other words, the pixel to which a vector **R** points, where **R** is a vector displacement in the image plane from p^s_{ij} calculated such that $|\mathbf{R}| = |\mathbf{g}^n(i,j)| = R$ and $\mathbf{R} \times \mathbf{g}(i,j) = 0$). The processor 120 determines which of the pixels in the transformed image data has the calculated offset **R** (rounded to the nearest pixel) in relation to the first selected pixel p^s_{ij} on the basis of the normalised pixel value gradient vector $\mathbf{g}^n(i,j)$ calculated for the first selected pixel p^s_{ij} in step S60. Unless the processor 120 employs a mask to exclude candidate locations for pixel p^o_{ij}, it may, as in the present embodiment, furthermore determine which of the pixels in the retinal image has an offset of **-R** (rounded to the nearest pixel) in relation to the first selected pixel p^s_{ij}, also on the basis of the normalised pixel value gradient vector $\mathbf{g}^n(i,j)$ calculated for the first selected pixel p^s_{ij} in step S60.

[0047] By the process in step S73, the processor 120 effectively registers a "vote" with each of two accumulators of the accumulator array A corresponding to respective pixels (at **R** and **-R** relative to p^o_{ij}) that are taken to be a candidates for the location of the fovea.

[0048] The process then proceeds to step S74, wherein the processor 120 determines whether the index I has reached the limit L, L being the number of pixels p^s_{ij} having been selected in step S50. If the counter I has not yet reached the limit L, then the counter I is incremented by 1 in step S75, and the process loops back to step S73, where the next selected pixel p^s_{ij} is processed as described above. In this way, each pixel in the image data at which the pixel value gradient magnitude exceeds the threshold contributes a vote in an accumulator corresponding to a pixel that is offset from the aforementioned pixel by a distance R, in a direction which is parallel to the local gradient and thus in a direction perpendicular to the local direction of the vascular structure.

[0049] Once all of the selected pixels p^s_{ij} have been processed (I = L in step S74), the process proceeds to step S76, wherein the processor 120 smooths the accumulated values in the accumulators to generate a response map. After a large number of selected pixels have been processed, a higher number of votes will have been registered in accumulators associated with pixels in the transformed image that are located near the centre of the aforementioned circle (along whose arcs the temporal vascular arcades extend), and this will be reflected in the response map. The array of accumulated values may be smoothed by applying any image-smoothing (or low-pass filtering) operation to the array. Examples of suitable image-smoothing methods are described in "Digital Image Processing" by R. C. Gonzalez and R. E. Woods (Pearson, August 31, 2007) and "Digital Image Processing and Computer Vision" by R. J. Schalkoff (John Wiley & Sons, September 2, 1992). The smoothing operation can be a linear operation such as convolution by a kernel which could have, for instance, a Gaussian shape or an array of similar values (to make a moving-average filter). Smoothing or low-pass filtering can also be performed in the frequency domain. The smoothing operation could also be a non-linear operation such as a median filter or a morphological filter. The processor 120 may, as in

the present embodiment, smooth the values in the accumulators using a Gaussian kernel which preferably has a standard deviation of 0.02 DFD to 0.5 DFD, and more preferably 0.075 DFD to 0.25 DFD, and yet more preferably 0.15 DFD. In any case, the response map is indicative of the probability distribution for the location of the fovea in the image of the retina.

[0050] Referring again to Fig. 4, after the completion of S70, the process proceeds to step S80, where the processor 120 estimates the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator of the array having accumulated an accumulated value which is a local or a global maximum among the accumulated values in the array of accumulators (the global maximum being the highest accumulated value of all the accumulators). The location of the fovea may be estimated as the location of the aforementioned pixel of the transformed image, or a pixel at the corresponding location in the received image (i.e. pre-transformation), or as the location of a pixel that is separated from either of these pixels by a predetermined displacement vector (determined, e.g. from experience or training). More generally, the processor 120 may estimate the location of the fovea using the location of a pixel in the received image data which is associated with an accumulator having accumulated an accumulated value that is within a predetermined range of values below a local or a global maximum among the accumulated values in the array of accumulators (the range, where it is non-zero, spanning, e.g. 5% or 10% of the maximum value). The locations of two or more local maxima determined from the response map may be used to estimate the location of the fovea. For example, multiple maxima may be chosen by taking all maxima which are above a certain threshold. Multiple maxima may be found, and a choice between which of these is associated with a pixel corresponding to the location of the fovea may be deferred to a subsequent data processing stage. A predetermined region of interest in the response map may be used to limit the choice of maxima. For example, the region of interest could be the region over which the fovea is known from experience to be very likely to occur.

[0051] In some embodiments (for example, where the accumulated values are negative or where the accumulators decrement an initially set count with each received "vote"), the processor 120 may estimate the location of the fovea in the retinal image as (or based on) the location of a pixel of the transformed image data which is associated with an accumulator of the array having accumulated an accumulated value which is a local or a global minimum among the accumulated values in the array of accumulators. Thus, the processor 120 may more generally determine the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator of the array having accumulated an accumulated value which is a (local or global) extremum of the accumulated values in the array of accumulators.

[Embodiment 2]

[0052] In the present embodiment, the processor 120 operates not on one but on N accumulator arrays, where $N \geq 2$. Each of these arrays of accumulators is associated with a

respective one of N different quantised directions along which the pixel value gradient vector **g** may be oriented, and is labelled with a corresponding index n, where n = 1, ... N. The possible directions along which **g** may be oriented are taken to fall within the range of angles [0, π], and this range is quantised into N sub-ranges or "quantised directions". For example, where N = 4, any vector **g**(i,j) may be classified as being oriented along a first quantised direction corresponding to angular range [0, π/4), a second quantised direction corresponding to angular range [π/4, π/2), a third quantised direction corresponding to angular range [π/2, 3π/4), or a fourth quantised direction corresponding to angular range [3π/4, π]. The function round[Nmod(θ,π)/π - 0.5], where θ = arctan($\mathbf{g}_y/\mathbf{g}_x$), $\mathbf{g}_y$ is the y-component of **g** and $\mathbf{g}_x$ is the x-component of **g,** may be used to quantise the direction of **g.** The accumulator array associated with the first quantised direction is assigned n = 1, the accumulator array associated with the second quantised direction is assigned n = 2, accumulator array associated with the third quantised direction is assigned n = 3, and accumulator array associated with the fourth quantised direction is assigned n = 4. In the present embodiment, an accumulator in the i^{th} column and the j^{th} row of the n^{th} accumulator array (associated with the n^{th} quantised direction) is denoted a_{ijn}. Each accumulator a_{ijn} in each accumulator array is associated with a respective pixel in the retinal image data, so that there are N accumulators (one in each of the N accumulator arrays) associated with each pixel in the transformed image data.

[0053] Figure 7 is a flow diagram illustrating the process by which the image-processing apparatus 170 of the present embodiment processes image data defining an image of a retina to determine the location of the fovea, which comprises steps S10 to S50 and the exemplary form of step S60 of the first embodiment described above (which is labelled S60' in Fig. 7), whose description will not be repeated here, as well as a sequence of operations performed on the N arrays of accumulators, which will now be described. It should be noted that the structure and operation of the present embodiment is the same as the first embodiment, apart from the differences hereinafter described. Furthermore, the possible modifications to the first embodiment mentioned above may also be made to the present embodiment.

[0054] In step S91, each accumulator of the N arrays of accumulators is initialised in the same way as in step S71 of Fig. 5, so that the accumulated value stored in each accumulator a_{ij} is zero.

[0055] In step S92, an index l, which is used to reference each of the selected pixels p^s_{ij} in turn, is set to an initial value of 1.

[0056] Then, in step S93, the processor 120 processes a first of the selected pixels p^s_{ij} (as l = 1 at this stage) by quantising the direction of either **g**(i,j) or $\mathbf{g}^n$(i,j) at the first selected pixel p^s_{ij} into one of the N quantised directions. Thus, one of the values of n (from 1 to N) may be associated with the first selected pixel p^s_{ij}, depending on the direction of **g** at that pixel, and this allows the corresponding accumulator array (having the same value of n) to be identified and accessed.

[0057] In step S94, the processor 120 adds a constant value (e.g. 1, as in the first embodiment) to an accumulator a_{ijn} of an accumulator array that is associated with the same quantised direction as the quantised direction of **g**(i,j) at the first selected pixel p^S_{ij}, wherein the accumulator to which the value is added is associated with a pixel p^O_{ij} of the transformed image data that is (as in the first embodiment) offset from the first selected pixel p^S_{ij} by the predetermined value R in a direction parallel to the direction of **g**(i,j) at the first selected pixel p^S_{ij}. Thus, in the present embodiment, which accumulator receives the "vote" for the selected pixel depends also on the direction of **g**(i,j) at the selected pixel, with the accumulator associated with the same quantised direction as the quantised direction of **g**(i,j), and which is associated with a pixel p^O_{ij} of the received image data that is offset from the selected pixel by distance R in a direction parallel to the direction of **g**(i,j) at the selected pixel, receiving the vote.

[0058] The process then proceeds to step S95, wherein the processor 120 determined whether the index I has reached the limit L, L being the number of pixels p^S_{ij} having been selected in step S50. If the counter I has not yet reached the limit L, then the counter I is incremented by 1 in step S96, and the process loops back to step S93, where the next selected pixel p^S_{ij} is processed as described above. In this way, each selected pixel in the transformed image data at which |**g**| exceeds the threshold contributes a vote in an accumulator a_{ijn} which is associated with the same quantised direction n as the quantised direction of **g** at the selected pixel, and which corresponds to a pixel that is offset from the selected pixel by a distance R, in a direction which is parallel to the local gradient at the selected pixel and thus in a direction normal to the local direction of the blood vessel at the selected pixel p^S_{ij}.

[0059] Once all of the selected pixels p^S_{ij} have been processed (I = L in step S95), the process proceeds to step S97, wherein the processor 120 smooths the accumulated values in the plurality of arrays of accumulators. The smoothing operation may be performed on accumulated values within one or more of the arrays, using the techniques described in connection with the first embodiment. Thus, Gaussian blurring of one or more accumulator arrays may be performed. Additionally or alternatively, the smoothing operation may be performed on values accumulated in the N accumulators a_{ijn} having the same indices i and j but different indices n (from 1 to N), using the same techniques.

[0060] In step S100, the processor 120 calculates, for each set of N accumulators a_{ijn} that are provided in different respective accumulator arrays and associated with the same pixel p^O_{ij} of the transformed image data, a product of the respective accumulated values that have been accumulated in the N accumulators a_{ijn}. The product is effectively a "soft AND" operation

across the different gradient orientations, so that the pixels in the image associated with high values of this product tend to be at locations about which blood vessels circle in the transformed image of the retina. In other words, wherever this product is non-zero, there must have been a contribution from each of the N directions associated with the N accumulator arrays.

[0061] The number of arrays of accumulators, N, is preferably such that $2 \leq N \leq 32$, and more preferably $4 \leq N \leq 16$. If N = 1 (as in the first embodiment) then the method is simplified in that there is no need to quantise the gradient orientation and there is no need for the "soft AND" combination of accumulator arrays. However, this simplification comes at a cost of a decrease in "signal-to-noise ratio" by a factor of 2 as compared to the case where N = 8, for example. Using values of N that are too high, on the other hand, increase the computational cost and may not allow the location of the optic disc or other anatomical feature to be determined reliably, as the product calculated in step S100 may be zero for all pixel locations. For the detection of the optic disc, it is preferable that N = 8. For the detection of vessel junctions, a smaller value of N may be used, for example 3 or 4.

[0062] The process proceeds from step S100 to step S110, wherein the processor 120 estimates the location of fovea in the image of the retina using the location of a pixel of the transformed image that is associated with a set of N accumulators, for which set the calculated product of accumulated values is a local or a global maximum among the calculated products of accumulated values. The location of the fovea may be estimated as the location of the aforementioned pixel of the transformed image, or a pixel at the corresponding location in the received image (i.e. pre-transformation), or as the location of a pixel that is offset in relation to either of these pixels by a predetermined amount (determined, e.g. from experience or training). More generally, the processor 120 may estimate the location of the fovea using the location of a pixel of the transformed image data that is associated with a set of N accumulators, for which set the calculated product of accumulated values is within a predetermined range of values below a (local or global) maximum among the calculated products of accumulated values. The locations of two or more local maxima determined from the response map may be used to estimate the location of the fovea. For example, multiple maxima may be chosen by taking all maxima which are above a certain threshold. Multiple maxima may be found, and a choice between which of these is associated with a pixel corresponding to the location of the fovea may be deferred. A predetermined region of interest in the response map may be used to limit the choice of maxima. For example, the region of interest could be the region over which the fovea is known from experience to be very likely to occur.

[Modifications and Variations]

[0063] Many further modifications and variations can be made to the embodiments described above.

[0064] For example, the first embodiment described above with reference to Figs. 2 to 6 may, as illustrated in Fig. 8, be modified to include an additional step after step S60' and before step S72, namely an additional step S120, wherein the processor 120 calculates a second normalised pixel value gradient vector $\mathbf{g}^n{}_2(i,j)$ for each of the selected pixels $p^s{}_{ij}$ by normalising the pixel value gradient vector **g** calculated for each of the selected pixels $p^s{}_{ij}$ so that $|\mathbf{g}^n{}_2(i,j)|$ at each of the selected pixels $p^s{}_{ij}$ takes a second common predetermined value R_2. The first embodiment may, as also illustrated in Fig. 8, be further modified to include an additional step S130 that is performed for each value of index I, wherein the processor 120 adds, for the I^{th} selected pixel, a respective value to the accumulator which is associated with a pixel in the transformed image data that is offset from the I^{th} selected pixel by the second predetermined value R_2 in the direction parallel to the direction of the pixel value gradient vector **g** at the I^{th} selected pixel. In other respects, this variant is the same as the first embodiment. The addition in step S130 may alternatively be made to a second accumulator layer that is associated with the second predetermined value R_2, with the corresponding elements in the two accumulator arrays being combined after step S74, by some hard or soft OR operation such as element-by-element addition.

[0065] More generally, more than two normalised pixel value gradient vectors may be calculated prior to step S72, preferably each having a magnitude within the preferred ranges of $|\mathbf{g}^n(i,j)|$ set out above. In these cases, the values may be accumulated in a single accumulator array, or in a corresponding number of accumulator arrays before being combined in the manner set out above. Using more than one calculated normalised pixel value gradient vector in the process may allow the location of the fovea to be determined more reliably.

[0066] The above-described modifications may also be made to the second embodiment described above.

[0067] As a further modification, which may be made to either of the above-described embodiments, or combined with any of the above-described modifications to these embodiments, the selection of pixels in step S50 may be omitted so that the processor 120 calculates, each of a plurality of pixels p_{ij} of the transformed (and optically preprocessed) image data, a respective pixel value gradient vector $\mathbf{g}(i,j)$ at the pixel. In this variant, the processor 120 would then calculate $\mathbf{g}^n(i,j)$ for each of the plurality of pixels p_{ij} by normalising the pixel value gradient vector calculated for each of the pixels p_{ij} so that the $|\mathbf{g}^n(i,j)|$ at each of the pixels takes a common predetermined value R. The processor 120 would operate on the array of accumulators by adding, for each of the plurality of pixels p_{ij}, a respective value to an accumulator associated with a pixel $p^o{}_{ij}$ of the transformed image data that is offset from the pixel of the plurality of pixels p_{ij} by the predetermined value R in a direction parallel to the direction of $\mathbf{g}(i,j)$ at the pixel of the plurality of pixels, and smoothing the values in the accumulators. In this variant, the processor may add, as the respective value, a respective

weighting to the accumulator associated with the pixel p^{o}_{ij}, the weighting being indicative of the magnitude of **g**(i,j) at the pixel of the plurality of pixels. The processor would then estimate the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator having accumulated an accumulated value which is within a predetermined range of an extremum of the accumulated values in the accumulators, as described in more detail above. In this variant, there is a contribution from each of the plurality of pixels to the accumulator, although the size of this contribution from each pixel will depend on the magnitude of **g**(i,j) at that pixel.

[0068] Furthermore, the ordering of some of the method steps in the embodiments and modifications thereof described above may be varied. For example, the ordering of step S60' and S120, and similarly S73 and S130, in Fig. 8 may be reversed, or these steps may be performed in parallel. Furthermore, the initialisation of accumulator array A may be performed at any point in the process before step S73 in the process illustrated in Fig. 5, at any point in the process before step S93 in the process illustrated in Fig. 7, and at any point before step S73 (step S130, where the order of S73 and S130 is reversed) in the process illustrated in Fig. 8 (the initialisation of course being outside the loop over index l).

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Non-patent literature cited in the description

- **A. D. FLEMING et al.**Automatic detection of retinal anatomy to assist diabetic retinopathy screeningPhys. Med. Biol., 2007, vol. 52, 331-345 [0003]
- **G. L. SHAPIRO et al.**Computer Vision - Chapter 10 Image Segmentation, 2000, [0004]
- **H. JELINEKM. J. CREE**Automated Image Detection of Retinal PathologyCRC Press20090101 [0023]
- **R. C. GONZALEZR. E. WOODS**Digital Image ProcessingPearson, 2007, [0036]
- **R. J. SCHALKOFF**Digital Image Processing and Computer VisionJohn Wiley & Sons19920902 [0036] [0049]
- Digital Image Processing**R. C. GONZALEZR. E. WOODS**Pearson20070831 [0049]

Patentkrav

1. Ikke-transient computerlæsbart lagringsmedium, der lagrer computerprogram-instruktioner, der, når de udføres af en processor, får processoren til at behandle billeddata, der definerer et billede af en vaskulær struktur af temporale vaskulære buegange af en nethinde for at bedømme en placering af nethindens fovea i billedet ved:

at modtage billeddataet;

at transformere det modtagne billeddata således at den vaskulære struktur i billedet defineret af det transformerede billeddata er mere cirkelformet end den vaskulære struktur i billedet defineret af det modtagne billeddata;

at beregne, for hver af en flerhed af pixler af det transformerede billeddata, en respektiv lokal orienteringsvektor som indikation af orienteringen af ethvert blodkar, som er til stede i billedet ved eller tilstødende pixelen;

at beregnet en normaliseret lokal orienteringsvektor for hver af flerheden af pixler ved at normalisere den lokale orienteringsvektor beregnet for hver af pixlerne således at størrelsen af den normaliserede lokale orienteringsvektor ved hver af pixlerne antager en fælles forudbestemt værdi;

at arbejde på en række af akkumulatorer, hvor hver akkumulator i rækken er associeret med en respektiv pixel af det transformerede billede, ved:

for hver af flerheden af pixler, at tilføje en respektiv værdi til en akkumulator associeret med en pixel af det transformerede billeddata, der er forskudt fra pixlen af flerheden af pixler med den forudbestemte værdi i en forudbestemt retning i forhold til retningen af af den lokale orienteringsvektor ved pixlen af flerheden af pixler; og

at udjævne værdierne i akkumulatorerne; og

at bedømme placeringen af fovea i billedet af nethinden under anvendelse af placeringen af en pixel af det transformerede billeddata, der er associeret med en akkumulator, der har akkumuleret en akkumuleret værdi, der er inden for et forudbestemt område af et maksimum af de akkumulerede værdier i akkumulatorerne.

2. Ikke-transient computerlæsbart lagringsmedium ifølge krav 1, hvor instruktionerne, når de udføres af processoren, får processoren til at udføre, forud for transformering af det modtagne billeddata, yderligere processer til:

(i) at behandle det modtagne billeddata for at generere billeddata med lavere rumlig opløsning som definerer et billede med lavere rumlig opløsning af den vaskulære struktur af de temporale vaskulære buegange på nethinden;

(ii) at generere filtreret billeddata ved at filtrere billeddata med den lavere rumlige opløsning for at forbedre lineære strukturer i billedet defineret af billeddata med lavere rumlig opløsning,

hvor processoren transformerer det filtrerede billeddata således at den vaskulære struktur i billedet defineret af det transformerede billeddata er mere cirkelformet end den vaskulære struktur i billedet defineret af det filtrerede billeddata.

3. Ikke-transient computerlæsbart lagringsmedium ifølge krav 2, hvor instruktionerne, når de udføres af processoren, får processoren til at udføre, mindst en gang, forud for transformering af det modtagne billeddata, yderligere processer til:

(iii) at behandle det filtrerede billeddata for at generere filtrerede billeddata med lavere rumlig opløsning der definerere et billede med lav rumlig opløsning af den vaskulære struktur af de temporale vaskulære buegange på nethinden;

(iv) at filtrere det filtrerede billeddata med lavere rumlig opløsning for at forbedre lineære strukturer i billedet defineret af det filtrerede billeddata med lavere rumlig opløsning,

hvor det filtrerede billeddata behandlet i enhver gentagelse af processen (iii) er det filtrerede billeddata genereret i den foregående udførelse af proces (iv), og hvor det filtrerede billeddata transformeret af processoren er det filtrerede billeddata genereret i den sidste udførelse af proces (iv).

4. Ikke-transient computerlæsbart lagringsmedium ifølge krav 3, hvor instruktionerne, når de udføres af processoren, får processoren til at reducere det modtagne billeddata med en faktor på 1,4 til 4 i udførelsen af proces (i), og til at reducere det filtrerede billeddata med en faktor på 1,4 til 4 i hver udførelse af

proces (iii).

5. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, får processoren til at arbejde på en flerhed, N, af rækker af akkumulatorer, hvor N er et heltal, som er lig med mindst 2, og hvor:

hver af rækkerne af akkumulatorer er associeret med en respektiv en af N kvantiserede retninger for den lokale orienteringsvektor;

hver af akkumulatorerne i hver af de N rækker er associeret med en respektiv pixel i det transformerede billeddata;

processoren arbejder på flerheden af rækker af akkumulatorer ved:

for hver af flerheden af pixler:

at kvantisere retningen af den beregnede lokale orienteringsvektor ved pixlen til en af de N kvantiserede retninger; og

at tilføje en respektiv værdi til en akkumulator i en akkumulatorrække, der er associeret med den samme kvantiserede retning som den kvantiserede retning af den lokale orienteringsvektor ved pixlen, idet akkumulatoren er associeret med en pixel af det transformerede billeddata, der er forskudt fra pixlen af flerheden af pixler med den forudbestemte værdi i den forudbestemte retning i forhold til retningen af den lokale orienteringsvektor ved pixlen af flerheden af pixler; og

at udjævne værdierne i akkumulatorerne af flerheden af rækker;

processoren beregner yderligere, for hvert sæt af N akkumulatorer, der er tilvejebragt i forskellige respektive akkumulatorrækker og associeret med den samme pixel af det transformerede billeddata, et produkt af de respektive akkumulerede værdier, der er blevet akkumuleret i N akkumulatorerne; og

processoren bedømmer placeringen af fovea i billedet af nethinden under anvendelse af placeringen af en pixel af det transformerede billeddata associeret med et sæt af N akkumulatorer, for hvilket sæt det beregnede produkt af akkumulerede værdier er inden for et forudbestemt område af et maksimum af de beregnede produkter af akkumulerede værdier.

6. Ikke-transient computerlæsbart lagringsmedium ifølge krav 5, hvor $2 \leq N \leq 32$, og mere fortrinsvis $4 \leq N \leq 16$.

7. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, yderligere får processoren til at beregne en anden normaliseret lokal orienteringsvektor for hver af flerheden af pixler ved at normalisere den lokale orienteringsvektor beregnet for hver af flerheden af pixler, således at størrelsen af den anden normaliserede lokale orienteringsvektor ved hver af pixlerne antager en anden fælles forudbestemt værdi, og til yderligere at arbejde på rækken af akkumulatorer ved:
for hver pixel af flerheden af pixler, at tilføje en respektiv værdi til en akkumulator associeret med en pixel i det transformerede billeddata, der er forskudt fra pixlen af flerheden af pixler med den anden forudbestemte værdi i den forudbestemte retning i forhold til retningen af pixelværdigradientvektoren ved pixlen af flerheden af pixler.

8. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, får processoren til at behandle billeddataet, der definerer billedet af den vaskulære struktur, for at bedømme en placering af fovea på nethinden i billedet ved at udføre en yderligere proces til:
at vælge pixler fra flerheden af pixler således at størrelsen af den beregnede lokale orienteringsvektor ved hver af de valgte pixler overstiger en tærskel, og ved:

at beregne den normaliserede lokale orienteringsvektor for hver af de valgte pixler ved at normalisere den lokale orienteringsvektor beregnet for hver af de valgte pixler, således at størrelsen af den normaliserede lokale orienteringsvektor ved hver af de valgte pixler antager en fælles forudbestemt værdi; og
at arbejde på rækken af akkumulatorer ved:
for hver af de valgte pixler, at tilføje en respektiv værdi til en akkumulator associeret med en pixel af det transformerede billeddata, der er forskudt fra den valgte pixel i billedet med den forudbestemte værdi i den forudbestemte retning i forhold til retningen af den lokale

orienteringsvektor ved den valgte pixel.

9. Ikke-transient computerlæsbart lagringsmedium ifølge krav 8, hvor instruktionerne, når de udføres af processor, får processoren til at vælge pixlerne fra flerheden af pixler, således at størrelsen af den beregnede lokale orienteringsvektor ved hver af de valgte pixler overstiger den 50. percentil, og mere fortrinsvis den 80. percentil, af lokale orienteringsvektorstørrelser af pixlerne i det transformerede billeddata.

10. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, får processoren til at arbejde på hver række af akkumulatorer ved:
for hver pixel af flerheden af pixler af det transformerede billeddata, at tilføje, som den respektive værdi, en respektiv vægtning af akkumulatoren associeret med en pixel i det transformerede billeddata, der er forskudt fra pixlen af flerheden af pixler med den forudbestemte værdi i den forudbestemte retning i forhold til retningen af den lokale orienteringsvektor ved pixlen af flerheden af pixler, idet vægtningen er indikation af størrelsen af den lokale orienteringsvektor ved pixlen af flerheden af pixler.

11. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, får processoren til at beregne den normaliserede lokale orienteringsvektor for hver pixel af flerheden af pixler, således at størrelsen af den normaliserede lokale orienteringsvektor ved hver pixel af flerheden af pixler er på mellem 0,5 DFD og 1,3 DFD, hvor DFD er afstanden mellem fovea og midten af en optisk plade i billedet defineret af det modtagne billeddata.

12. Ikke-transient computerlæsbart lagringsmedium ifølge et hvilket som helst foregående krav, hvor instruktionerne, når de udføres af processoren, får processoren til at udjævne værdierne i akkumulatorerne under anvendelse af en kerne, som har en standardafvigelse på 0,02 DFD til 0,5 DFD, og mere fortrinsvis 0,075 DFD til 0,25 DFD, hvor DFD er afstanden mellem fovea og midten af en optisk plade i billedet defineret af det modtagne billeddata.

13. Billedbehandlingsapparat omfattende en processor og en hukommelse, idet hukommelsen lagrer instruktioner, som kan udføres af processoren, hvorved processoren kan fungere til at behandle billeddata, der definerer et billede af en vaskulær struktur af temporale vaskulære buegange på en nethinde for at bedømme en placering af fovea på nethinden i billedet ved:

at modtage billeddataet;

at transformere det modtagne billeddata således at den vaskulære struktur i billedet defineret af det transformerede billeddata er mere cirkelformet end den vaskulære struktur i billedet defineret af det modtagne billeddata;

at beregne, for hver af en flerhed af pixler af det transformerede billeddata, en respektiv lokal orienteringsvektor som indikation af retningen af ethvert blodkar, som er til stede i billedet ved eller tilstødende pixlen;

at beregne en normaliseret lokal orienteringsvektor for hver af flerheden af pixler ved at normalisere den lokale orienteringsvektor beregnet for hver af pixlerne, således at størrelsen af den normaliserede lokale orienteringsvektor ved hver af pixlerne antager en fælles forudbestemt værdi;

at arbejde på en række af akkumulatorer, hvor hver akkumulator i rækken er associeret med en respektiv pixel af det transformerede billeddata, ved:

for hver af flerheden af pixler, at tilføje en respektiv værdi til en akkumulator associeret med en pixel af det transformerede billeddata, der er forskudt fra pixlen af flerheden af pixler med den forudbestemte værdi i en forudbestemt retning i forhold til retningen af den lokale orienteringsvektor ved pixlen af flerheden af pixler; og

at udjævne værdier i akkumulatorerne; og

at bedømme placeringen af fovea i billedet af nethinden under anvendelse af placeringen af en pixel af det transformerede billeddata, der er associeret med en akkumulator, der har akkumuleret en akkumuleret værdi, der er inden for et forudbestemt område af et maksimum af de akkumulerede værdier i akkumulatorerne.

14. Billedbehandlingsapparat ifølge krav 13, hvor instruktionerne, når de udføres af processoren, får processoren til at udføre, forud for transformering af det modtagne billeddata, yderligere processer til:

(i) at behandle det modtagne billeddata for at generere billeddata med lavere rumlig opløsning, der definerer et billede med lavere rumlig opløsning af den vaskulære struktur af de temporale vaskulære buegange på nethinden;

(ii) at generere filtreret billeddata ved at filtrere billeddataet med lavere rumlig opløsing for at forøge lineære strukturer i billedet defineret af billeddataet med lavere rumlig opløsning,

hvor processoren transformerer det filtrerede billeddata således at den vaskulære struktur i billedet defineret af det transformerede billeddata er mere cirkelformet end den vaskulære struktur i billedet defineret af det modtagne billeddata.

15. Billedbehandlingsapparat ifølge krav 14, hvor instruktionerne, når de udføres af processoren, får processoren til at udføre, mindst en gang, forud for transformering af det modtagne billeddata, yderligere processer til:

(iii) at behandle det filtrerede billeddata for at generere filtrerede billeddata med lavere rumlig opløsning, der definerer et billede med lav rumlig opløsning af den vaskulære struktur af de temporale vaskulære buegange på nethinden;

(iv) at filtrere de filtrerede billeddata med lavere rumlig opløsning for at forøge lineære strukturer i billedet defineret af de filtrerede billeddata med lavere rumlig opløsning,

hvor det filtrerede billeddata behandlet i enhver gentagelse af processen (iii) er det filtrerede billeddata genereret i den forudgående udførelse af proces (iv), og hvor det filtrerede billeddata transformeret af processoren er det filtrerede billeddata genereret i den sidste udførelse af proces (iv).

DRAWINGS

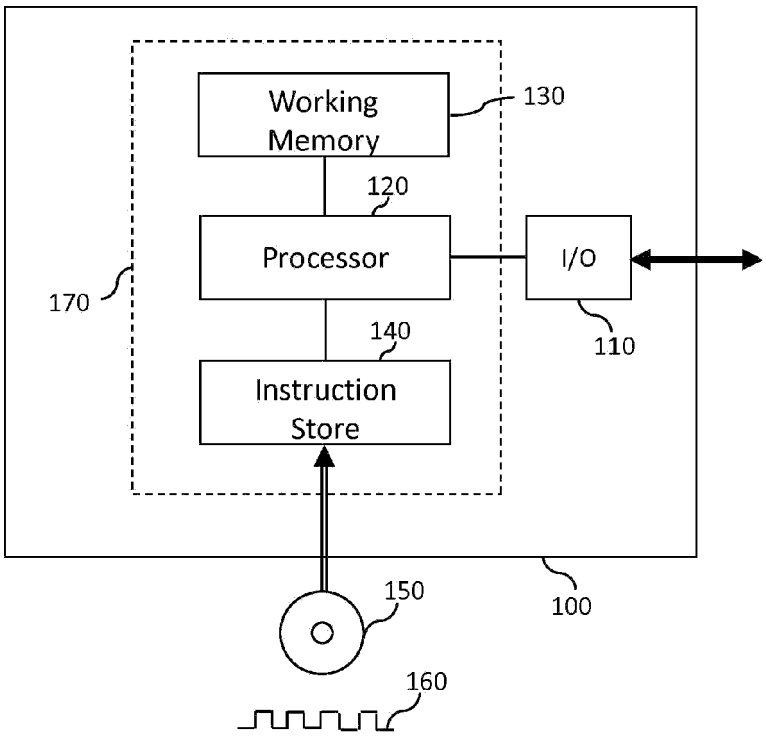


Fig. 1

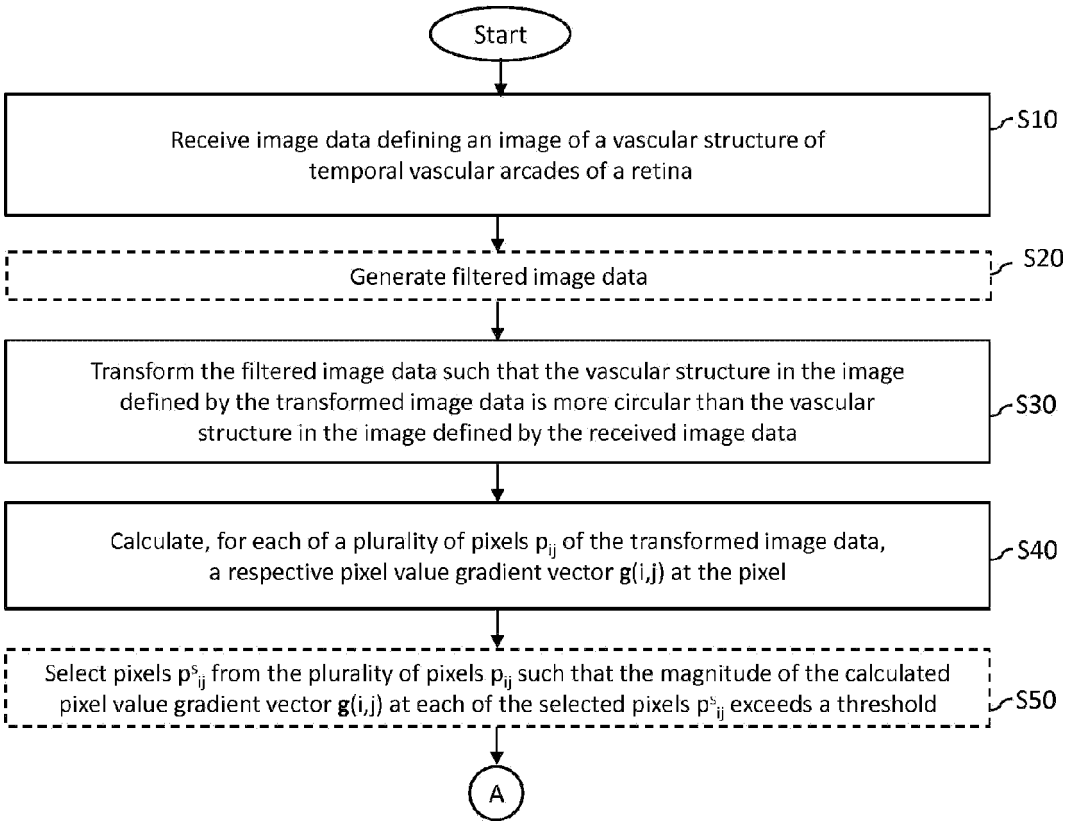

Start
Receive image data defining an image of a vascular structure of temporal vascular arcades of a retina
S10
Generate filtered image data
S20
Transform the filtered image data such that the vascular structure in the image defined by the transformed image data is more circular than the vascular structure in the image defined by the received image data
S30
Calculate, for each of a plurality of pixels p_{ij} of the transformed image data, a respective pixel value gradient vector $\mathbf{g}(i,j)$ at the pixel
S40
Select pixels p^s_{ij} from the plurality of pixels p_{ij} such that the magnitude of the calculated pixel value gradient vector $\mathbf{g}(i,j)$ at each of the selected pixels p^s_{ij} exceeds a threshold
S50
A

Fig. 2



Fig. 3A

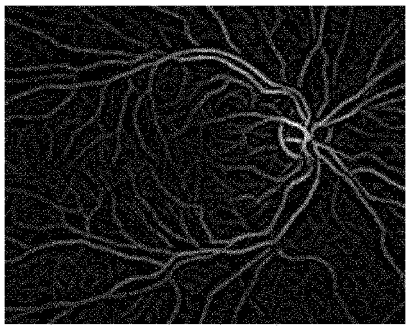


Fig. 3B

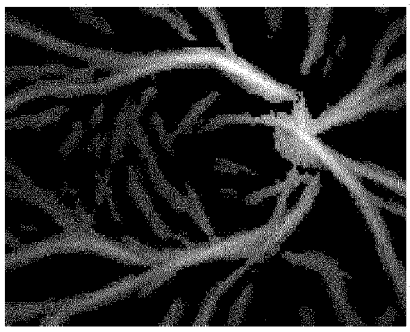


Fig. 3C

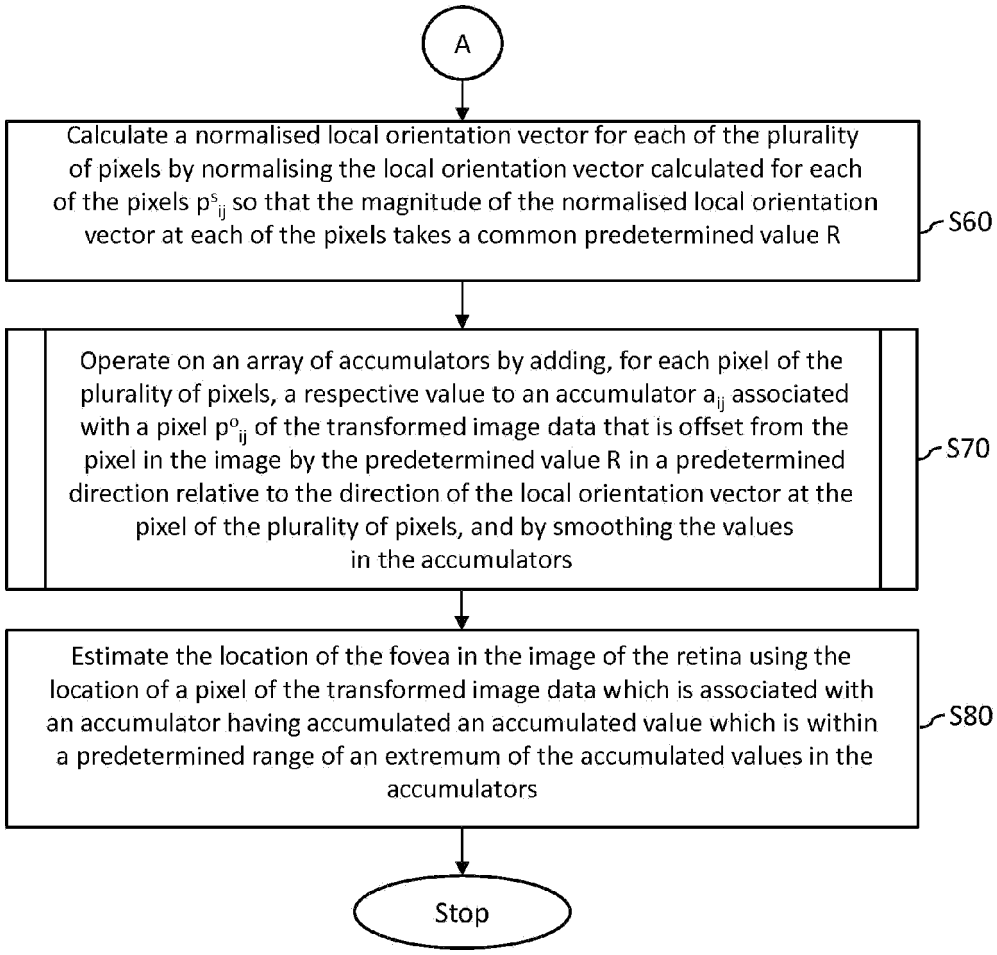

A
Calculate a normalised local orientation vector for each of the plurality of pixels by normalising the local orientation vector calculated for each of the pixels p^s_{ij} so that the magnitude of the normalised local orientation vector at each of the pixels takes a common predetermined value R
S60
Operate on an array of accumulators by adding, for each pixel of the plurality of pixels, a respective value to an accumulator a_{ij} associated with a pixel p^o_{ij} of the transformed image data that is offset from the pixel in the image by the predetermined value R in a predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, and by smoothing the values in the accumulators
S70
Estimate the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator having accumulated an accumulated value which is within a predetermined range of an extremum of the accumulated values in the accumulators
S80
Stop

Fig. 4

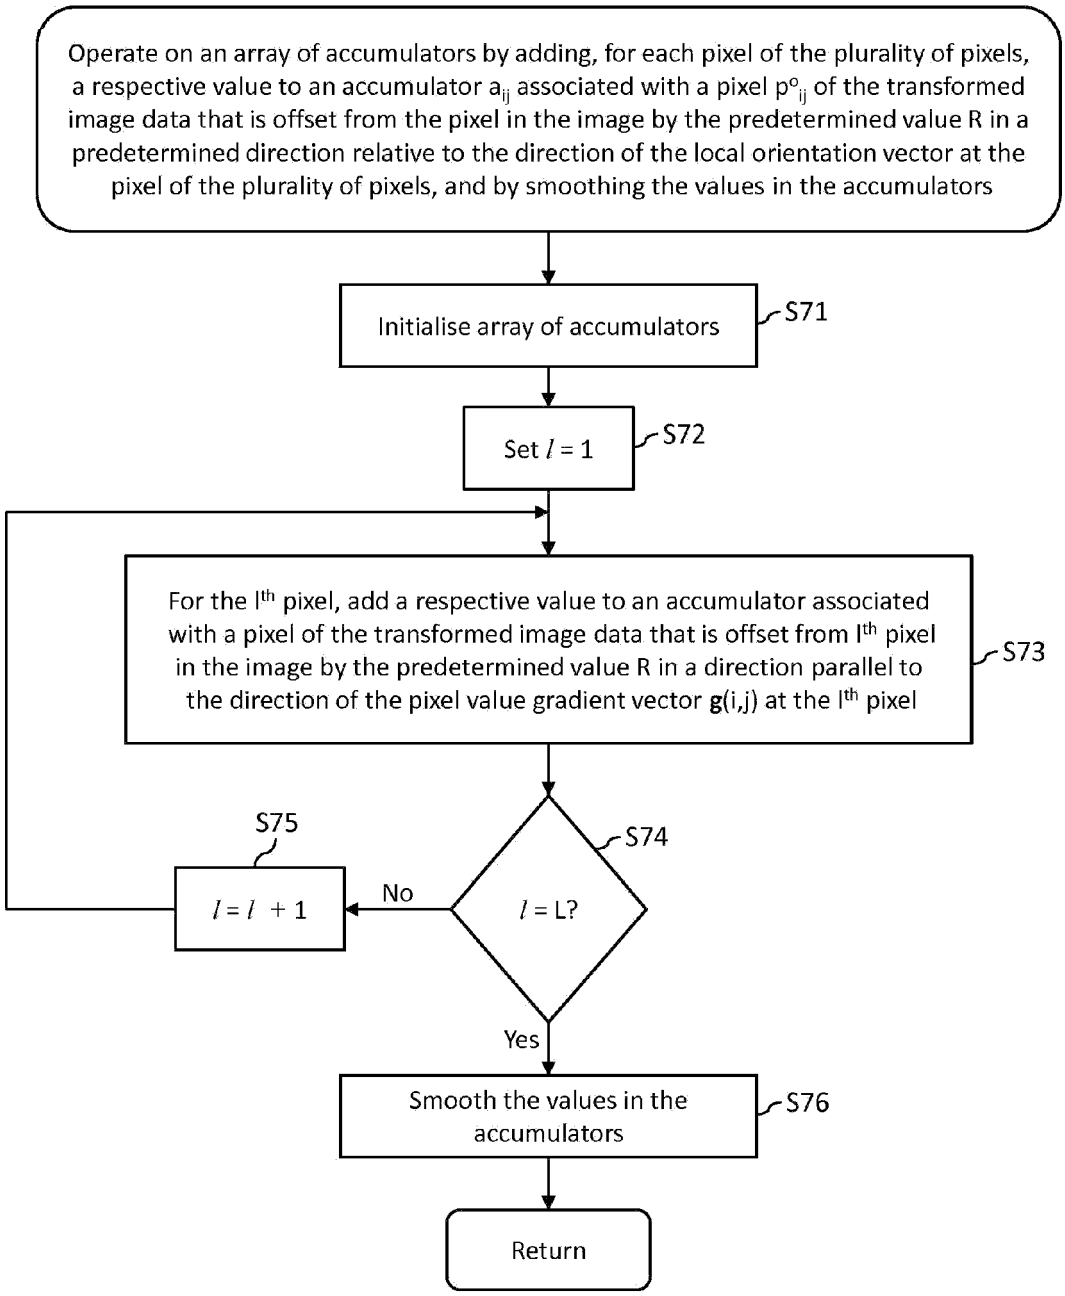

Operate on an array of accumulators by adding, for each pixel of the plurality of pixels, a respective value to an accumulator a_{ij} associated with a pixel p^o_{ij} of the transformed image data that is offset from the pixel in the image by the predetermined value R in a predetermined direction relative to the direction of the local orientation vector at the pixel of the plurality of pixels, and by smoothing the values in the accumulators
Initialise array of accumulators
S71
Set l = 1
S72
For the l^{th} pixel, add a respective value to an accumulator associated with a pixel of the transformed image data that is offset from l^{th} pixel in the image by the predetermined value R in a direction parallel to the direction of the pixel value gradient vector g(i,j) at the l^{th} pixel
S73
S74
l = L?
No
S75
l = l + 1
Yes
Smooth the values in the accumulators
S76
Return

Fig. 5

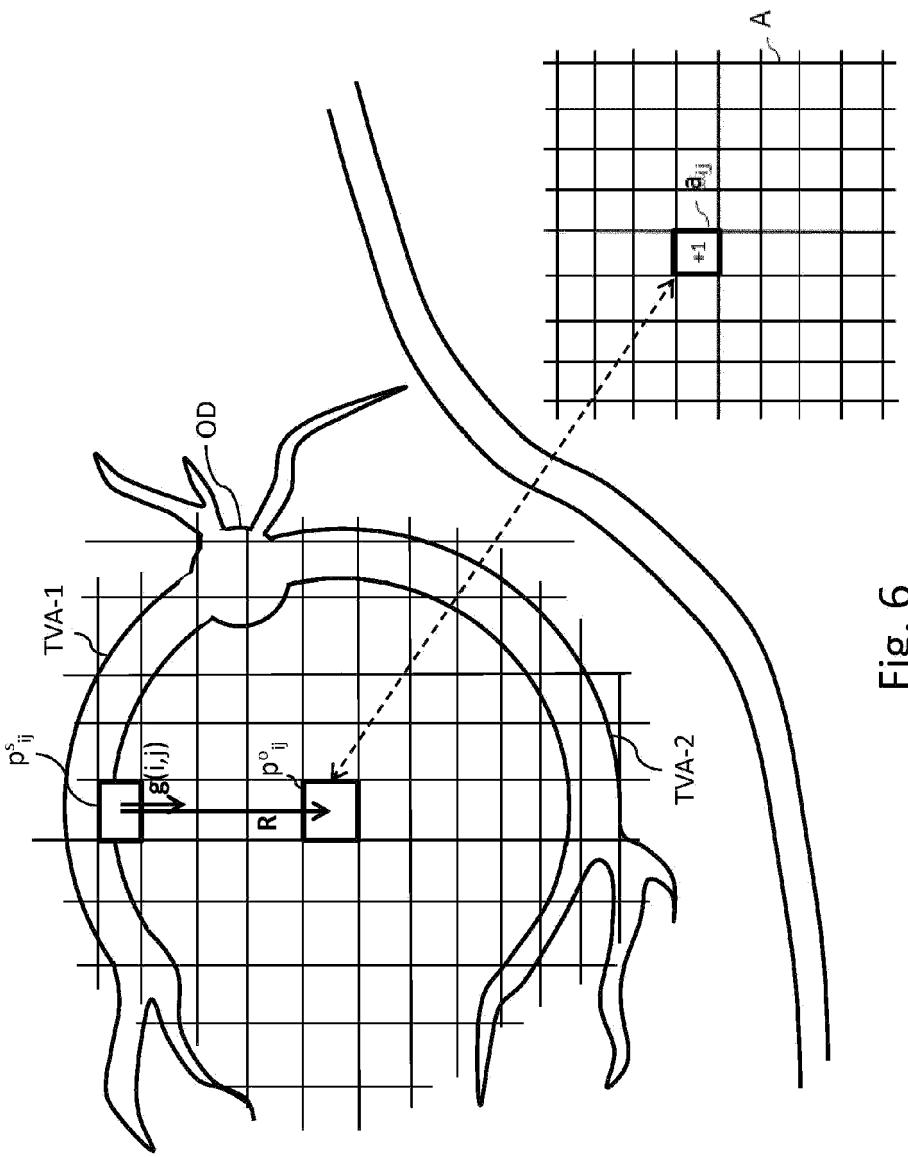


Fig. 6

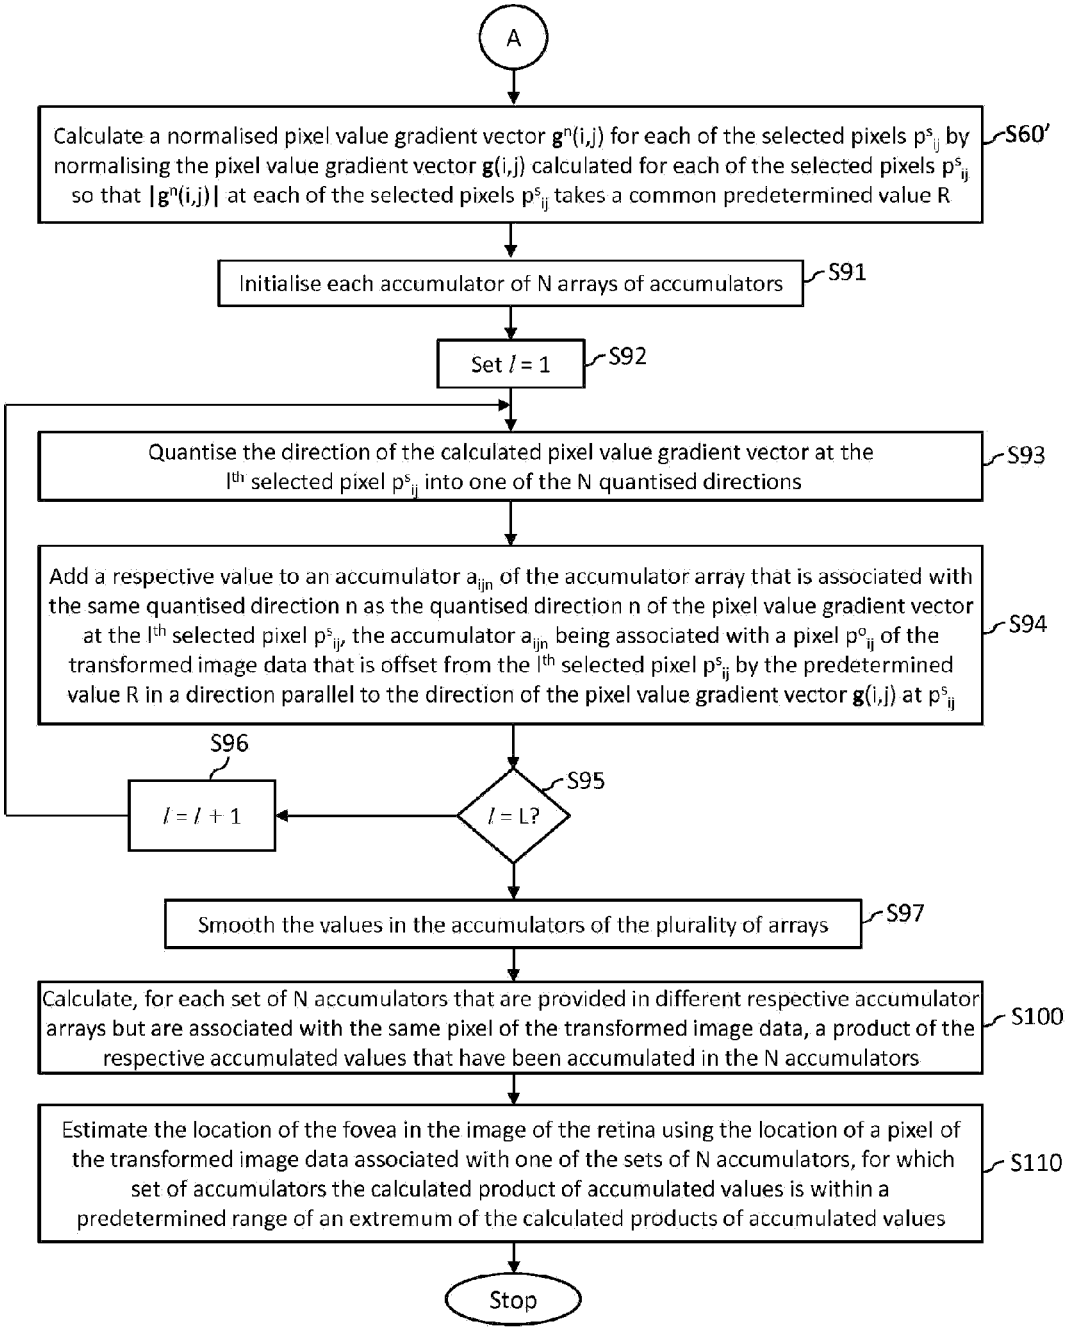

A
Calculate a normalised pixel value gradient vector $\mathbf{g}^n(i,j)$ for each of the selected pixels p^s_{ij} by normalising the pixel value gradient vector $\mathbf{g}(i,j)$ calculated for each of the selected pixels p^s_{ij} so that $|\mathbf{g}^n(i,j)|$ at each of the selected pixels p^s_{ij} takes a common predetermined value R
S60'
Initialise each accumulator of N arrays of accumulators
S91
Set l = 1
S92
Quantise the direction of the calculated pixel value gradient vector at the l^{th} selected pixel p^s_{ij} into one of the N quantised directions
S93
Add a respective value to an accumulator a_{ijn} of the accumulator array that is associated with the same quantised direction n as the quantised direction n of the pixel value gradient vector at the l^{th} selected pixel p^s_{ij}, the accumulator a_{ijn} being associated with a pixel p^o_{ij} of the transformed image data that is offset from the l^{th} selected pixel p^s_{ij} by the predetermined value R in a direction parallel to the direction of the pixel value gradient vector $\mathbf{g}(i,j)$ at p^s_{ij}
S94
S96
$l = l + 1$
S95
l = L?
Smooth the values in the accumulators of the plurality of arrays
S97
Calculate, for each set of N accumulators that are provided in different respective accumulator arrays but are associated with the same pixel of the transformed image data, a product of the respective accumulated values that have been accumulated in the N accumulators
S100
Estimate the location of the fovea in the image of the retina using the location of a pixel of the transformed image data associated with one of the sets of N accumulators, for which set of accumulators the calculated product of accumulated values is within a predetermined range of an extremum of the calculated products of accumulated values
S110
Stop

Fig. 7

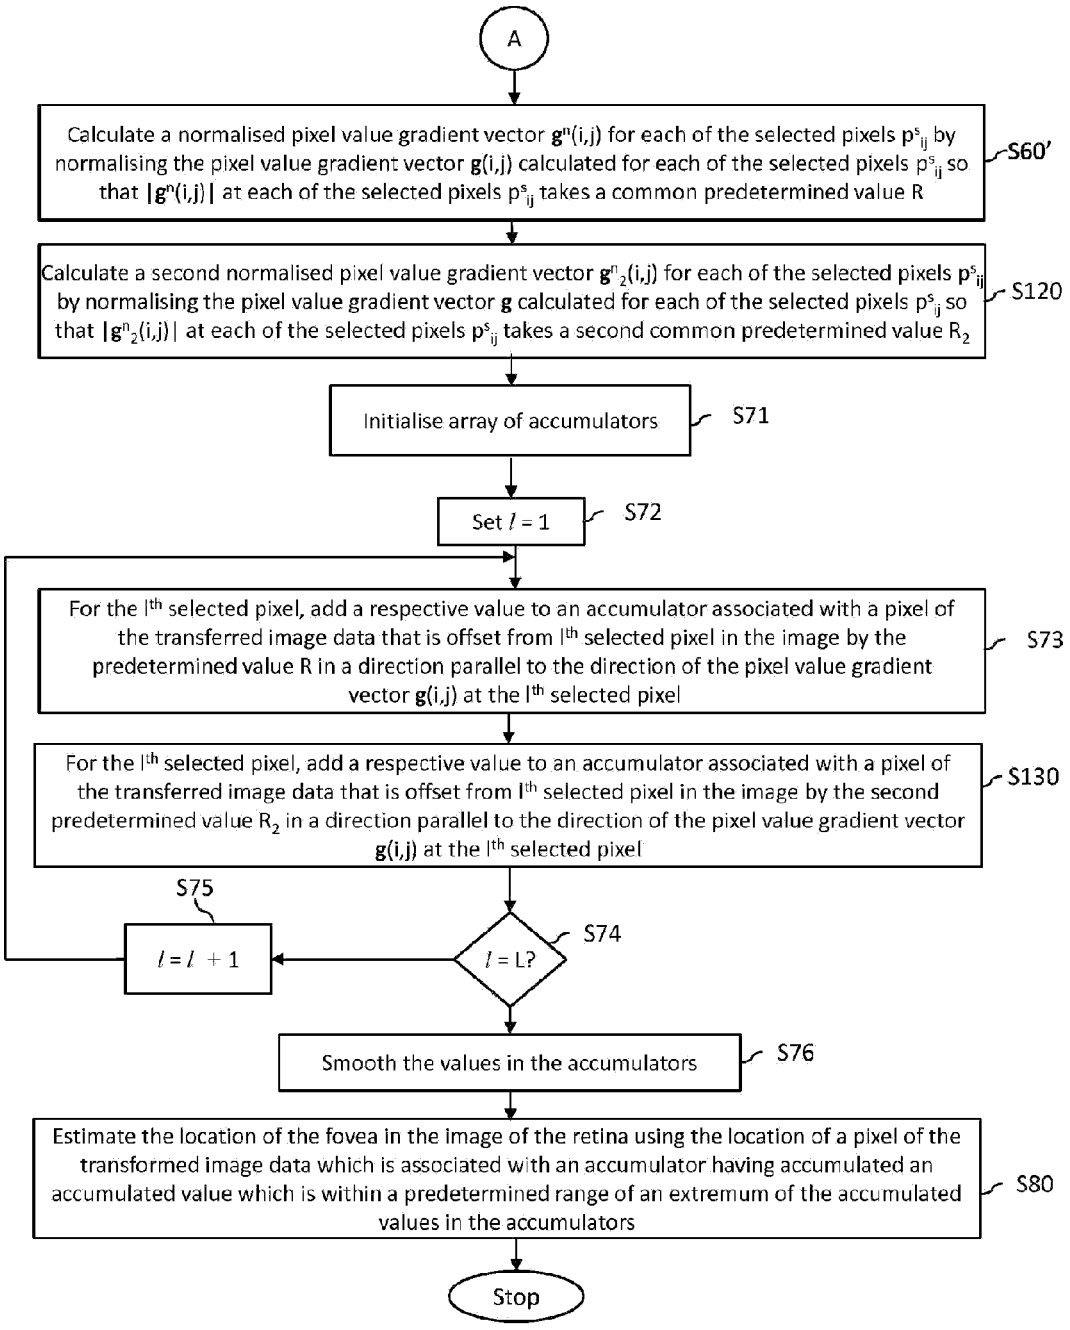

A
Calculate a normalised pixel value gradient vector $\mathbf{g}^n(i,j)$ for each of the selected pixels p^s_{ij} by normalising the pixel value gradient vector $\mathbf{g}(i,j)$ calculated for each of the selected pixels p^s_{ij} so that $|\mathbf{g}^n(i,j)|$ at each of the selected pixels p^s_{ij} takes a common predetermined value R
S60'
Calculate a second normalised pixel value gradient vector $\mathbf{g}^n{}_2(i,j)$ for each of the selected pixels p^s_{ij} by normalising the pixel value gradient vector $\mathbf{g}$ calculated for each of the selected pixels p^s_{ij} so that $|\mathbf{g}^n{}_2(i,j)|$ at each of the selected pixels p^s_{ij} takes a second common predetermined value R_2
S120
Initialise array of accumulators
S71
Set l = 1
S72
For the l^{th} selected pixel, add a respective value to an accumulator associated with a pixel of the transferred image data that is offset from l^{th} selected pixel in the image by the predetermined value R in a direction parallel to the direction of the pixel value gradient vector $\mathbf{g}(i,j)$ at the l^{th} selected pixel
S73
For the l^{th} selected pixel, add a respective value to an accumulator associated with a pixel of the transferred image data that is offset from l^{th} selected pixel in the image by the second predetermined value R_2 in a direction parallel to the direction of the pixel value gradient vector $\mathbf{g}(i,j)$ at the l^{th} selected pixel
S130
S75
$l = l + 1$
l = L?
S74
Smooth the values in the accumulators
S76
Estimate the location of the fovea in the image of the retina using the location of a pixel of the transformed image data which is associated with an accumulator having accumulated an accumulated value which is within a predetermined range of an extremum of the accumulated values in the accumulators
S80
Stop

Fig. 8