



PATENT SPECIFICATION

(11) S84986

(21) Application No. S2008/0340

(22) Date of Filing of Application: 30/04/2008

(30) Priority Data:

(31) 11/937,377 (32) 08/11/2007 (33) United States of America (US)

(45) Specification Published: 01 October 2008

(51) Int. Cl. (2006)
 G06K 9/00
 G06T 7/00

(54) Title: Detecting red-eye defects in digital images

(72) Inventor: MIHAI CIUC

(73) Patent Granted to: **FOTONATION VISION LIMITED, An Irish Company, Galway Business Park, Dangan, Galway, Ireland**

Detecting Redeye Defects in Digital Images

The present invention relates to a method and apparatus for detecting redev
defects in digital images. In the present specification a redev defect means any
5 flash-induced artefact in an image of a human or animal eye, whether actually
red or not.

BACKGROUND OF THE INVENTION

- 10 Redeye is a phenomenon in flash photography where a flash is reflected within a
subject's eye and appears in a photograph as a light, usually red, dot where the black
pupil of the subject's eye would normally appear. The unnatural glowing red of an
eye is due to internal reflections from the vascular membrane behind the retina,
which is rich in blood vessels. This objectionable phenomenon is well understood to
15 be caused in part by a small angle between the flash of the camera and the lens of
the camera. This angle has decreased with the miniaturization of cameras with
integral flash capabilities. Additional contributors include the relative closeness of
the subject to the camera and ambient light levels.
- 20 The redev phenomenon can be reduced by causing the iris to reduce the opening of
the pupil. This is typically done with a "pre-flash", a flash or illumination of light
shortly before a flash photograph is taken. This causes the iris to close.
Unfortunately, the pre-flash is an objectionable 0.2 to 0.6 seconds prior to the flash
photograph. This delay is readily discernible and easily within the reaction time of a
25 human subject. Consequently the subject may believe the pre-flash is the actual
photograph and be in a less than desirable position at the time of the actual
photograph. Alternately, the subject must be informed of the pre-flash, typically
loosing any spontaneity of the subject captured in the photograph.
- 30 Digital photography eliminates the need for film as the image is digitally captured
and stored in a memory array for display on a display screen on the camera itself.
This allows photographs to be viewed and enjoyed virtually instantaneously as
opposed to waiting for film processing. Digital photography equipment includes
microprocessors for image processing and compression and camera systems control.

It is possible to exploit the computation capabilities of such microprocessors for performing operations to improve redevye detection and elimination.

Existing techniques for redevye detection and correction in digital images are
5 described in US 6,407,777 (DeLuca) and US 2005/0232490. However, these prior
methods are not efficient in practice.

The object of the present invention is to mitigate the disadvantages associated with
the prior art.

10

DISCLOSURE OF THE INVENTION

The present invention provides a method for detecting a redevye defect in a digital
image as claimed in claim 1.

15

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention will now be described by way of example with
reference to the accompanying drawings, in which:

20 **Figure 1** is a block diagram of a digital camera operating in accordance with an
embodiment of the present invention.

Figure 2 is a flow diagram of the steps performed by software in the camera of **Figure**
1 of an embodiment of redevye detection and correction according to the invention.

25

Figures 3 to 7 show image data processed according to the method of **Figure 2** (for
clarity, **Figures 6(a) to 6(d), 6(f)** and **7** are the negative or inverted form of the
relevant images described in the text).

30 DESCRIPTION OF THE PREFERRED EMBODIMENT

In the present specification, the term "image" refers to image data and does not
necessarily imply that an actual viewable image is present at any stage of the
processing.

FIG. 1 is a block diagram of a digital image acquisition device 20 which in the present embodiment is a portable digital camera, and includes a processor 120. It can be appreciated that many of the processes implemented in the digital camera may be implemented in or controlled by software operating in a microprocessor, central processing unit, controller, digital signal processor and/or an application specific integrated circuit, collectively depicted as processor 120. Generically, all user interface and control of peripheral components such as buttons and display is controlled by a microcontroller 122. The processor 120, in response to a user input at 122, such as half pressing a shutter button (pre-capture mode 32), initiates and controls the digital photographic process. Ambient light exposure is monitored using light sensor 40 in order to automatically determine if a flash is to be used. A distance to the subject is determined using a focus component 50 which also focuses the image on image capture component 60. If a flash is to be used, processor 120 causes the flash 70 to generate a photographic flash in substantial coincidence with the recording of the image by image capture component 60 upon full depression of the shutter button. The image capture component 60 digitally records the image in colour. The image capture component preferably includes a CCD (charge coupled device) or CMOS to facilitate digital recording. The flash may be selectively generated either in response to the light sensor 40 or a manual input 72 from the user of the camera. The high resolution image recorded by image capture component 60 is stored in an image store 80 which may comprise computer memory such a dynamic random access memory or a non-volatile memory. The camera is equipped with a display 100, such as an LCD, for preview and post-view of images.

In the case of preview images which are generated in the pre-capture mode 32 with the shutter button half-pressed, the display 100 can assist the user in composing the image, as well as being used to determine focusing and exposure. Temporary storage 82 is used to store one or more of the preview images and can be part of the image store 80 or a separate component. The preview image is preferably generated by the image capture component 60. For speed and memory efficiency reasons, preview images preferably have a lower pixel resolution than the main image taken when the shutter button is fully depressed, and are generated by sub-sampling a raw captured image using software 124 which can be part of the general processor 120 or dedicated hardware or combination thereof. Depending on the settings of this

hardware subsystem, the pre-acquisition image processing may satisfy some predetermined test criteria prior to storing a preview image. Such test criteria may be chronological, such as to constantly replace the previous saved preview image with a new captured preview image every 0.5 seconds during the pre-capture mode
5 32, until the high resolution main image is captured by full depression of the shutter button. More sophisticated criteria may involve analysis of the preview image content, for example, testing the image for changes, before deciding whether the new preview image should replace a previously saved image. Other criteria may be based on image analysis such as the sharpness, or metadata analysis such as the
10 exposure condition, whether a flash is going to happen, and/or the distance to the subject.

If test criteria are not met, the camera continues by capturing the next preview image without saving the current one. The process continues until the final high
15 resolution main image is acquired and saved by fully depressing the shutter button.

Where multiple preview images can be saved, a new preview image will be placed on a chronological First In First Out (FIFO) stack, until the user takes the final picture. The reason for storing multiple preview images is that the last preview image, or any
20 single preview image, may not be the best reference image for comparison with the final high resolution image in, for example, a red-eye correction process or, in the present embodiment, mid-shot mode processing. By storing multiple images, a better reference image can be achieved, and a closer alignment between the preview and the final captured image can be achieved in an alignment stage
25 discussed later.

The camera is also able to capture and store in the temporary storage 82 one or more low resolution post-view images. Post-view images are low resolution images essentially the same as preview images, except that they occur after the main high
30 resolution image is captured.

A redeye detection and correction filter 90 can be integral to the camera 20 or part of an external processing device 10 such as a desktop computer, a colour printer or a photo kiosk. In this embodiment, the filter 90 receives the captured high resolution
35 digital image from the store 80 and analyzes it, 92, to detect redeyes. The analysis

92 is performed according to the principles of the invention as described in the embodiments to follow. If redeves are found, the filter modifies the image, 94, to remove the redevye from the image using well-known techniques. The modified image may be either displayed on image display 100, saved on a persistent storage 5 112 which can be internal or a removable storage such as CF card, SD card or the like, or downloaded to another device via image output means 110 which can be tethered or wireless. The redevye filter 90 can be brought into operation either automatically each time the flash is used, or upon user demand via input 30. Although illustrated as a separate item, where the filter 90 is part of the camera it 10 may be implemented by suitable software on the processor 120.

Figure 2 is a flow diagram of the steps performed by software in the camera and/or external processing device of Figure 1 of an embodiment of redevye detection and correction according to the invention.

15

1. Eye detection

The first step 200 of the embodiment is to run an eye detection algorithm to identify regions of the image suspected to contain an eye. Examples of such algorithms are 20 described in PCT Application No. PCT/EP2006/008342 (Ref: FN128) and US Application Nos. 60/865,375 filed November 10, 2006; US 60 865,622 filed November 13, 2006; and US 60/915,669 filed May 2, 2007 (Ref: FN181/FN208) whose output is a detection rectangle which ideally should contain an eye but might be a false positive (a non-eye region). Moreover, even in the case where the detection rectangle 25 contains an eye, in most cases it is a non-defect eye, which should not be modified by a subsequent correction algorithm. Therefore, the detection rectangle should be filtered to discard both non-eye and normal-eye regions before applying correction.

The main difficulty of the detection task is the wide variety of redevyes, both in terms 30 of color (covering all possible hues from bright red to orange, yellow, white, and also combinations of them) and size. Moreover, in many cases, the defect is far from being uniform in terms of color and/or intensity. Thus, the object is to find an invariant to all of the defect eyes, i.e., a property that holds true in substantially all cases. This property is that the defect is brighter than its neighborhood, and is 35 surrounded, in general, by a darker ring.

Thus, the detection of a defect is performed solely on the intensity image (computed in this embodiment as the average of the red and green components), whereas color information is used at the later stage to make decisions upon correcting or not a defect region.

2. Detection of a defect

The next step 202 therefore comprises computing the intensity image, Figure 3a, of the detection rectangle issued by the eye detection algorithm at step 200.

In order to detect the defect within the intensity image, the present method exploits the fact that it is brighter than its close neighborhood. However, attempting to separate the defect by a mere thresholding of the intensity image has little chance of succeeding. This is because, firstly, the range of luminosities of possible defects is wide, therefore no prior threshold can be set to work in all cases. Also, using adaptive thresholds, i.e. determining a threshold for each case based on local intensity information, would not work well in general. This can be explained as follows: in many cases, the border between the defect and its surrounding (iris, eyelids) is not well defined all around the defect. Figure 3a presents an example of this situation that occurs quite often - the intensity transition between the defect and the upper eyelid is not very prominent, especially in comparison with the transition between the defect and the iris. This is seen in Figure 3b which is an intensity profile along the vertical line seen in Figure 3a, the two points P1, P2 marking the transition intensities between the defect and upper eyelid and iris respectively. In order to spatially separate the defect by simple thresholding, a large threshold would have to be used - larger than 160, according to P1 in Figure 3b - which would not allow inclusion in the detected region of all pixels belonging to the defect. As one can see in Figure 3c, which is a reduced intensity image of the defect thresholded with the lowest threshold that spatially isolates the defect, there is a bright ring that is not caught in the detected region owing to using too large a threshold.

Therefore, in order to be able to separate the defect from surrounding, a more complex technique has been devised. It is based on the observation that, in most of the cases, no matter the intensity at which the defect/surrounding transition occurs, the transition is valley-shaped.

5

2.1 Mean-shift eye segmentation.

Accordingly, at step 204, the intensity image of Figure 3a is subjected to mean-shift segmentation. The mean-shift procedure (D. Comaniciu, P. Meer: Mean Shift: A Robust Approach toward Feature Space Analysis, IEEE Trans. Pattern Analysis Machine Intell., Vol. 24, No. 5, 603-619, 2002) has been used for a number of different tasks, the most important of which is clustering (segmentation). Given an N-dimensional distribution, the basic idea of the algorithm is to identify the modes of the distribution (a mode being defined by a local maximum). As an illustration, Figure 4 presents an example of mode identification of a 1D distribution, the short vertical lines identifying the transitions from one mode to the next. The approach is a bottom-up one, in that it starts from the points with lowest values and builds the mode evolving towards local maxima. In order to identify the modes, the following simple procedure is carried out:

20

- For each point of the distribution, a new mode is initialized. Then, the neighbor with the maximum gradient is sought and added to the current mode.
- If the maximum gradient is positive, the procedure continues in the same manner starting from the newly-added neighbor.
- If the maximum gradient is negative (i.e., the current point is a local maximum) the procedure stops.
- If, during the procedure, the neighbor characterized by the maximum gradient had already been inspected (that is, it has already been assigned to a mode previously), then all points belonging to the current mode are assigned to the mode of the winning neighbor.
- The procedure stops when all the points of the distribution have been assigned to a mode.

25

30

This procedure is generally applied to a distribution (a probability density function, namely, a histogram). In step 204, it is applied directly on the intensity image, considering it as a surface (equivalent to a 2D distribution). If the intensity image is too large, a local averaging with a small kernel (5x5 at most) may be applied prior to
5 extracting local modes, in order to reduce small variations that may cause over-segmentation (detection of too many modes).

The results of mean-shift intensity segmentation are shown in Figure 5, of which Figure 5a is a pseudo 3D mesh of the intensity image with local maxima extracted by
10 the mean shift segmentation, Figure 5b shows the regions resulting from the segmentation, and Figure 5c shows the local maxima of the regions in Figure 5b viewed as an image. One may observe two things. First, as expected, the eye defect region is split into multiple segments, referred to as defect sub-segments, due to the fact that there are several local maxima inside the defect region - see Figure 3b. In
15 this case there are five defect sub-segments (labelled 1 to 5 in Figure 5b) that belong to the eye defect. However, the defect is well separated from its surrounding - the eyelid is separated from the defect even though the border between the two is very weak at some pixels.

20 2.2 Merging the defect sub-segments

Next, step 206, the defect sub-segments composing the eye defect are merged into a single segment.

25 The observation that enables merging all defect sub-segments into one segment is that, in general, the local maxima of these sub-segments are high and grouped together into a very high intensity region. Thus, this high-intensity region is extracted by thresholding with a very high threshold and, in one implementation, identifying all segments in the mean-shift segmented intensity that have their
30 maxima located inside this region. Another criterion for identifying segments is whether portion of a segment area, for example more than 75%, lies inside the region.

The algorithm has the following steps:

35

1. Automatically compute the high threshold by the following procedure:
 - a. Identify the highest-intensity pixel inside the intensity image of Figure 3a.
 - b. Take a smaller rectangle centered on that pixel. The size of the rectangle depends upon the size of the eye-detection rectangle: e.g. 10x10 for $H < 100$, 20x20 for $100 < H < 200$, 30x30 for $H > 200$, with H being the height of the eye-detection rectangle. In the example of Figure 3(a), the size H of the eye detection rectangle is 20x20.
 - c. Compute the average intensity of the pixels inside the smaller rectangle.
 - d. Assign the threshold to the sum of the average intensity above and a fixed quantity. The fixed quantity is determined heuristically, and the value used in the present embodiment is 25.
2. Threshold the intensity image with the above-determined threshold - the result is shown in Figure 6a.
3. Identify all connected components of the thresholded image - Figure 6b. This comprises linking all connected pixels in the thresholded image into respective groups 600a...f (the smallest groups are not labelled).
4. Eliminate connected components (i.e. the groups not labelled in Figure 6b) whose size is inappropriate (e.g. they are smaller than 0.1% of the eye-detection rectangle area or larger than 4% of the area) or that touch the border of the detection rectangle - Figure 6c.
5. Rank the remaining components with respect to their average intensity and retain only the top two.
6. If the difference between the average intensities of the two components is significant (the difference between them exceeds 10% of the highest average intensity), the component with the highest average intensity is selected, otherwise the component with the highest average saturation is selected - Figure 6d.
7. Identify all the segments in the mean-shift segmented image whose maxima are located inside the selected component (or who satisfy any other criterion relative to the selected component) - Figure 6e.
8. Merge all segments and discard all other pixels in the intensity image - Figure 6f.

As seen in Figure 6f, the region determined by this procedure contains not only the bright defect, but also some parts belonging to the surroundings that should not be

corrected (the whole being referred to hereinafter as the “extended” defect). Therefore, these parts are removed, step 208.

2.3 Extracting the final defect

5

In order to segment the real defect from the extended defect, step 208 applies a histogram thresholding technique to the part of the intensity image corresponding to the extended defect. The mean-shift segmentation step 204 ensures that the extended defect contains no other bright parts than the defect, and this is central for determining the threshold on the histogram. The technique is the following: the intensity histogram of the part of the intensity image corresponding to the extended defect is computed. As the size of the extended defect is, in general, small, the histogram is very likely to be noisy (i.e. to contain many “spikes” - see Figure 7a). Therefore, prior to computing the threshold, it is smoothed by convolving it with an n-pixel wide averaging kernel. The width n of the kernel is chosen heuristically and in the present embodiment n=31. Typically, after smoothing, the histogram is bimodal - see the smooth line in Figure 7a. The first prominent maximum is determined, then the prominent minimum that follows is taken as the sought threshold, indicated by the vertical line in Figure 7a. Note that the selected minimum is not the *absolute* minimum between the two modes, but it is a *local* minimum: it has (in this embodiment) 3 neighbours to the left and 3 to the right with higher values. Also, it is the first local minimum that follows the *first* maximum. All pixels in the extended defect having an intensity lower than the threshold are removed. Figure 7b shows the result, obtained by thresholding the image in Figure 6f with the threshold computed as above. What remains is potentially the real defect that must undergo correction. However, as previously stated, a decision on the likelihood that the detected region is a real defect must be made prior to applying correction.

30 3. Filtering

In the previous section, the procedure to identify the best candidate for the defect region inside the eye detection rectangle has been described. However, the detection rectangle might either be a false positive (i.e., a non-eye region wrongly detected by the eye detector) or contain a non-defect eye. In both cases, correction

35

should not be applied. Therefore, the intensity image of Figure 7b is filtered, step 210, to decide whether the defect region determined by the method described above is indeed an eye defect or not. The most likely defect candidate detected for a non-defect eye is the white part of the eye (sclera), as it is the brightest region in the detection rectangle. Correcting it would result in a very nasty false positive therefore it must not pass the filtering stage.

The principle behind the filter step 210 is that a real defect has the following properties:

- 10 - it is round;
 - it is more saturated than the rest of the detection rectangle;
 - it is yellower or redder (or both) than the rest of the detection rectangle;
 - its contour is darker than its inside.
- 15 Therefore, to decide whether an eye-defect candidate region is a real defect, or not, the following parameters are taken into account:
- The region's roundness (computed as the classical circularity factor, i.e. the squared perimeter to area ratio).
 - 20 - The average saturation values of the region and of the detection rectangle;
 - The average a values of the region and of the detection rectangle (standing for average degrees of redness).
 - The average b values of the region and of the detection rectangle (standing for average degrees of redness). Both a and b are coordinates of the representation of colors in Lab color space.
 - 25 - The ratio of the average intensity of the region's contour to the average intensity of the region (this measure is the only one computed on the extended defect and not on the final defect).

30 Based on these measures, step 210 includes a number of filter stages to assess if the general conditions appertaining to a defect are met. The filter stages take into account that few defects meet all of the four characteristics enumerated above. Therefore, each of the conditions above might be relaxed to a certain extent at the expense of the others (i.e. imposing more strict conditions on the other characteristics). For instance, an average saturation of the candidate region slightly

35

smaller than that of the detected rectangle is acceptable provided that the region is very round, and very yellow, etc.

4. Correction

5

If the defect candidate region has passed all filters and has been declared a defect, it must undergo correction, which should consist in reducing the intensity of the defect. Yet, a number of additional issues should be addressed, namely, the fact that the defect might contain a glint, and also the fact that reducing the intensity of pixels inside the defect might create unpleasant artefacts at the region's border.

10

4.1 Glint detection

Accordingly, a glint detection filter is applied at step 212. In general, glint is a region that is both brighter and less saturated than the rest of the defect. Hence, the glint detection filter applies the following steps:

15

- Detect the brightest $n\%$ points of the detected defect region.
- Detect the least saturated $m\%$ points of the detected defect region. Both m and n are chosen, as a function of the size of the defect region, based on the observation that the larger the defect area, the smaller the area occupied by the glint. Also, $n > m$, i.e. we allow inspection of more bright pixels than least saturated pixels because, in some defects, there are parts of the defect that are brighter than the glint (e.g. for defects having between 100 and 300 pixels, $m=15$, $n=30$, etc.).
- Intersect the two sets of points.
- Identify the connected components of the intersection.
- Choose the less saturated one as the candidate glint.
- If the chosen region is round enough (roundness being assessed in terms of both aspect ratio and filling factor) it is declared glint, therefore it is removed from the region to be corrected.

20

25

30

4.2 Correction of the defect

After glint detection, the region to be corrected is available. Correction then proceeds at step 214 by reducing the intensity of pixels inside the detected defect by a fixed factor, for example 5 (thus from 120 to 24, from 210 to 42, from 109 to 22, etc.). In order to reduce the artefacts at the defect border (artefacts created by the fact that pixels located on either side of the defect border are treated so much differently), after darkening the defect region, a 3x3 blurring is applied on the inner and outer borders of the defect. Also, another correction which is applied is redness correction in a small bounding box surrounding the defect. Its aim is to eliminate the reddish parts of the eye that do not get included in the defect region, parts that still remain visible in a number of cases unless they are processed in this manner.

The invention is not limited to the embodiment(s) described herein but can be amended or modified without departing from the scope of the present invention.

CLAIMS

1. A method for detecting a redeye defect in a digital image containing an eye, the method comprising the following steps:
 - 5 (a) converting the digital image into an intensity image,
 - (b) segmenting at least a portion of the intensity image into segments each having a local intensity maximum,
 - (c) thresholding a corresponding portion of the digital image to identify regions of relatively high intensity,
 - 10 (d) selecting a region from at least some of the regions of step (c) having substantially the highest average intensity, and
 - (e) identifying segments from step (b) intersecting the region selected at step (d) according to a pre-determined criterion.
- 15 2. The method of claim 1 further comprising identifying regions from step (c) having a size falling within a predetermined range, and wherein said region of substantially the highest average intensity is selected from said regions.
3. The method of claim 1 wherein step (e) comprises identifying those segments
20 from step (b) whose maxima are located in the region selected at step (d).
4. The method of claim 1, wherein the average intensities of the two regions from step (c) having the highest and next highest average intensities differ by more than a predetermined amount, and the region selected at step (d) is the
25 region having the highest average intensity.
5. The method of claim 1, wherein the average intensities of the two regions from step (c) having the highest and next highest average intensities differ by less than a predetermined amount, and the region selected at step (d) is the
30 region having the highest average saturation.

011103

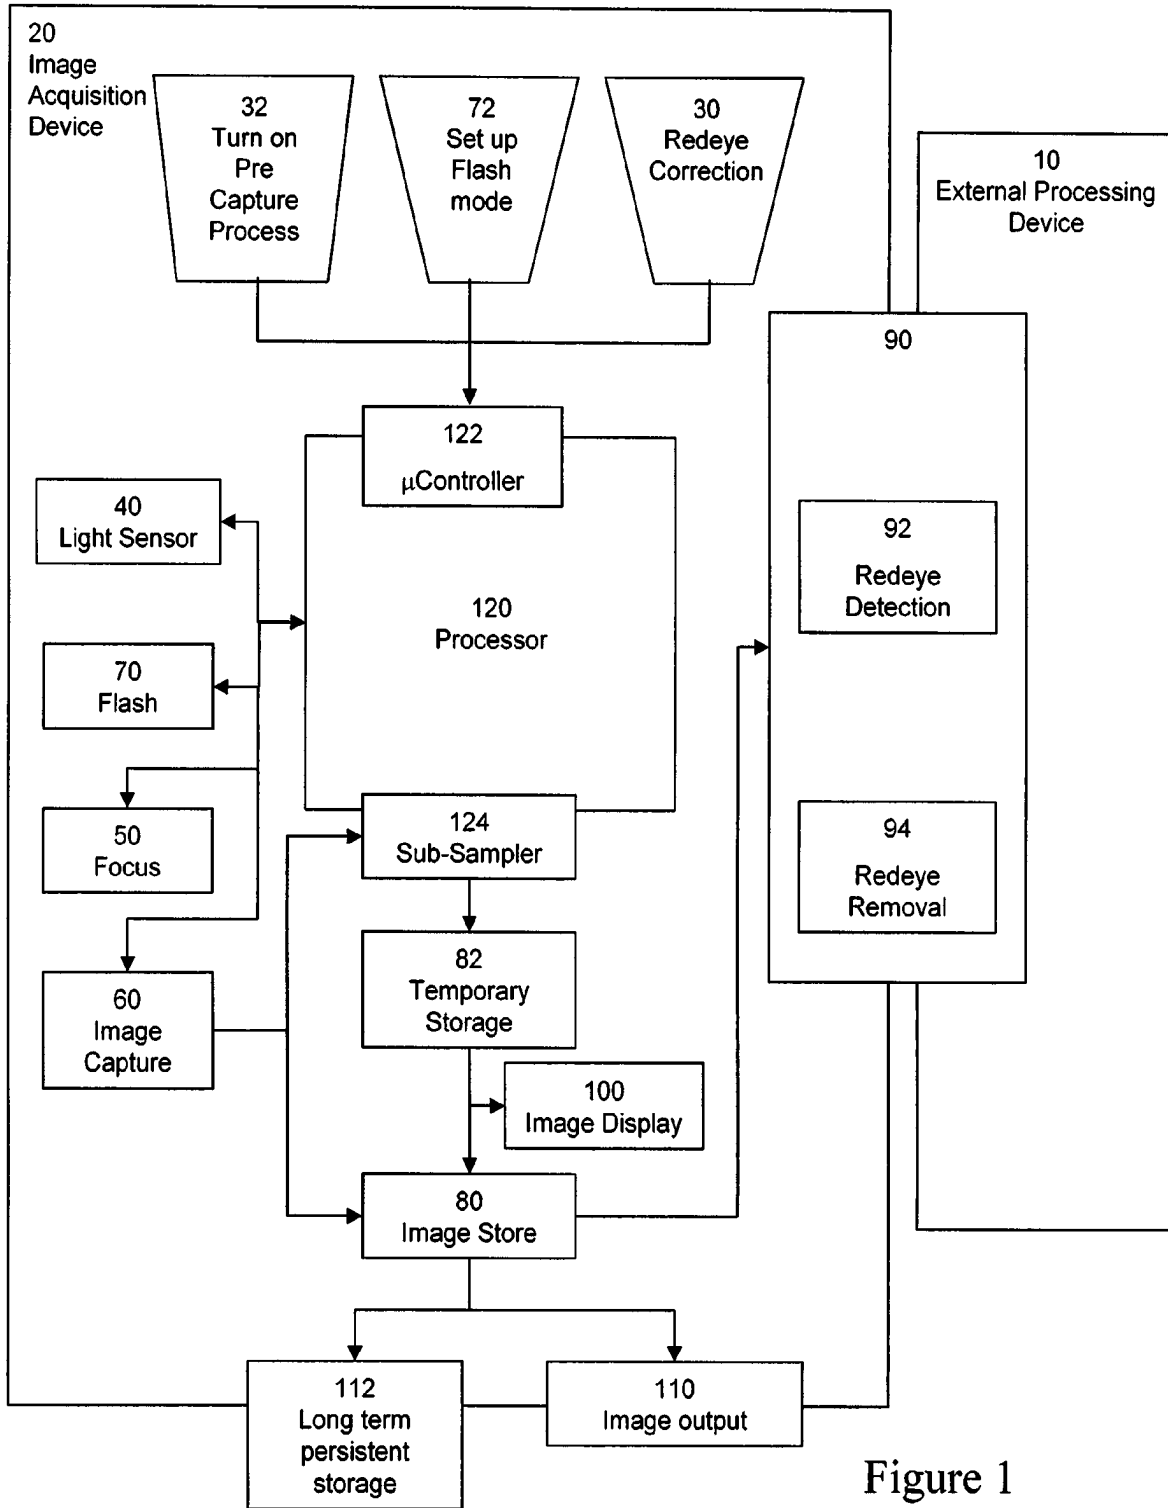


Figure 1

01-18-02

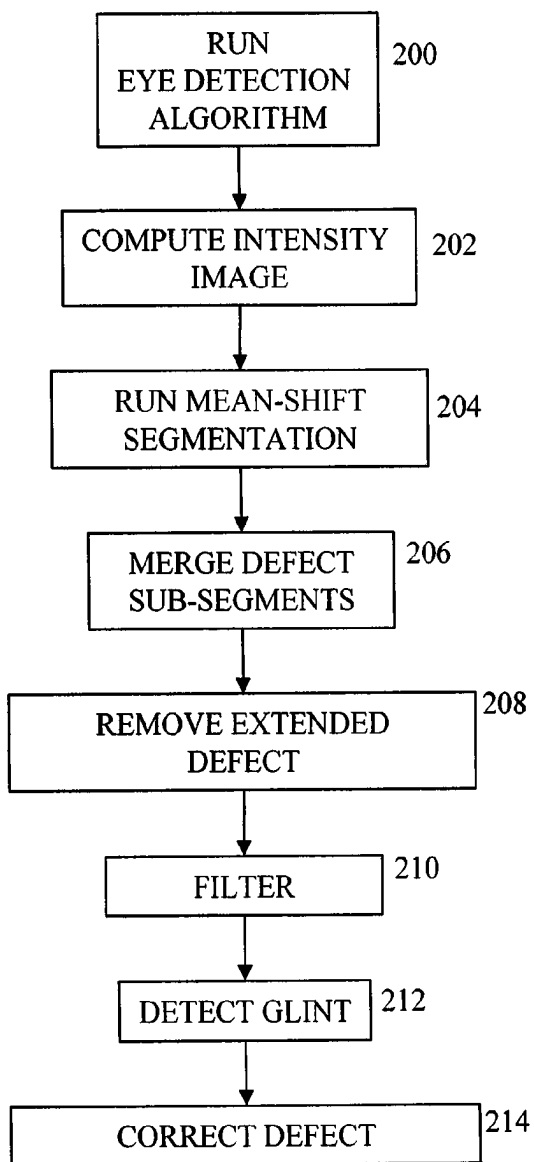





Figure 2



(a)



(c)



(b)

Figure 3




Figure 4









Figure 5




(a)




(b)




(c)



(d)





(e)



(f)

Figure 6

20000000



(b)

Figure 7