



- (51) International Patent Classification:
G06T 7/00 (2006.01)
- (21) International Application Number:
PCT/GB2015/051405
- (22) International Filing Date:
13 May 2015 (13.05.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
1408499.0 13 May 2014 (13.05.2014) GB
- (71) Applicant: **STRING LABS LIMITED** [GB/GB]; 194B Addington Road, Selsdon, South Croydon, Surrey CR2 8LD (GB).
- (72) Inventor: **ØVERBYE, Johan**; c/o String Labs Limited, 194B Addington Road, Selsdon, South Croydon, Surrey CR2 8LD (GB).
- (74) Agent: **SLINGSBY PARTNERS LLP**; 5 Chancery Lane, London, Greater London EC4A 1BL (GB).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: IDENTIFYING FEATURES




FIG. 1

(57) Abstract: A method for identifying features in an image, the method comprising: a) from a first location, scanning the image along a scan path in a scan direction away from the first location so as to detect a border in the image at a second location; b) mapping a path defined by the border in a first direction along the border from the second location, the path being mapped in dependence on a border quality measure; c) determining whether the mapped path forms a perimeter; and d) if the mapped path does form a perimeter, scanning the image from the mapped path so as to detect a further border in the image; or if the mapped path does not form a perimeter, mapping a path defined by the border from the second location in a second direction along the border that is different to the first direction.

IDENTIFYING FEATURES

This invention relates to detecting features in images.

It is commonplace for computing devices to gather, process and display images. For example, many mobile phones, tablet computers and personal computers can collect images from an on-board camera, process those images and display the processed images on a screen. The processing could involve a very simple algorithm such as converting a still image from colour to greyscale. More complex algorithms might include noise reduction and tonal enhancement, and might be applied to video rather than to still photos.

An important field of image analysis is border or border detection. Border detection involves estimating the location of borders between regions of an image. Those regions might correspond to objects as perceived by a viewer to be represented in the image. Once the locations of the borders have been estimated additional processing could be carried out on the image. In one example, the additional processing could involve highlighting a border in a bright colour in order to draw attention to an object. This might be useful in providing night vision capabilities for vehicles. In a second example, the additional processing could involve re-colouring the area of an image surrounded by a border in order to create a visual effect. In a third example, the additional processing could involve inferring a three-dimensional relationship between objects whose borders have been estimated, either from a single image or from a series of images of a similar scene taken at different locations.

Although border detection can provide many advantages, it is a complex problem, and is particularly difficult to implement on a device that has relatively little computing power. An example of this is when border detection is to be performed on a mobile phone. The computing power of mobile phones is increasing, but it is still very limited; and since mobile phones are battery powered it is desirable to minimise the amount of energy they expend in performing video processing, to thereby improve battery life.

Examples of border detection algorithms are disclosed in EP 2 615 597 A, US 2014/072208, WO 2013/186662 and CN 103440644 A.

In some applications it is desirable to be able to process video data in real time. Where the video is made up of a series of frames, real time processing involves processing a frame of the video in

less time – at least on average – than the period between successive frames. In one example application, a device has a camera and a display screen that can display the video stream being captured by the camera. It might be desired to manipulate the displayed video stream, for example to overlay on the captured video pre-prepared graphical representations which move on the display as the camera zooms and pans to give the illusion that they are present in real life.

It would be desirable to have an algorithm that minimises the amount of computing power required to detect objects and features in images. It would be desirable for such an algorithm to be implemented on a portable device such as a mobile phone to analyse frames in real-time video of reasonable resolution.

According to the present invention there is provided a method for identifying features in an image, the method comprising: a) from a first location, scanning the image along a scan path in a scan direction away from the first location so as to detect a border in the image at a second location; b) mapping a path defined by the border in a first direction along the border from the second location, the path being mapped in dependence on a border quality measure; c) determining whether the mapped path forms a perimeter; and d) if the mapped path does form a perimeter, scanning the image from the mapped path so as to detect a further border in the image; or if the mapped path does not form a perimeter, mapping a path defined by the border from the second location in a second direction along the border that is different to the first direction.

According to a second aspect of the present invention there is provided a device comprising: a memory storing non-transient program code; and a processor configured to identify features in an image by executing the code to perform the steps of: a) from a first location, scanning the image along a scan path in a scan direction away from the first location so as to detect a border in the image at a second location; b) mapping a path defined by the border in a first direction along the border from the second location, the path being mapped in dependence on a border quality measure; c) determining whether the mapped path forms a perimeter; and d) if the mapped path does form a perimeter, scanning the image from the mapped path so as to detect a further border in the image; or if the mapped path does not form a perimeter, mapping a path defined by the border from the second location in a second direction along the border that is different to the first direction

The method may comprise: on mapping said path in the second direction, scanning the image from the mapped path so as to detect a further border in the image.

The method may comprise: on detecting a further border in the image at a third location, repeating steps b) to d), wherein the repeated steps are performed from the third location.

The method may comprise, after having traced a border: determining whether the circumference or length of the border and/or the compounded quality measurements of the border exceeds a predetermined threshold; and if not, discarding the border.

The method may comprise, from a fourth location: scanning the image along a second scan path in the scan direction away from the fourth location so as to detect a border in the image at a fifth location, the fourth location being offset by a predetermined distance from the first location in a direction perpendicular to the scan direction.

The method may comprise, on detecting the border: ignoring the border if it has been mapped previously.

In each step of scanning the image, the scan direction may be the same.

The image may be composed of a plurality of pixels. The method may comprise mapping a path by detecting borders between pixel positions. The quality measure may be indicative of the quality of the border at each border position.

The method may comprise: estimating an attribute for each of a set of two adjacent pixel positions in the image; assessing whether a predetermined binary condition differs in respect of the two pixel positions, and if so, determining that a border is present in the part of the image represented by the two pixel positions.

The border quality measure may be a value derived from neighbouring pixels. The method may further comprise: determining whether the derived value is above or below a predetermined threshold, and if above, mapping that border position as part of the path; or if below, not mapping that border position as part of the path.

The attribute of the pixel position may be based on a Laplacian of Gaussian filter applied to the image and centred on that pixel position.

The attribute of a pixel position may be based on estimating the second-order derivative of an isotropic Gaussian convolution of the image at the pixel position and with respect to the gradient direction of the Gaussian.

A second attribute of a pixel position may be based on estimating the third-order derivative of an isotropic Gaussian convolution of the image at the pixel position and with respect to the gradient direction of the Gaussian. The border quality measure at a given border position may be an interpolation of the second attribute with respect to each of the two pixel positions on either side of the border position.

The image may be a frame of a video stream.

The device may comprise a camera. The image may be an image captured by the camera.

The image may be a frame of a video stream captured by the camera.

The device may be configured to perform the said steps in real time on successive frames of the video stream.

The present invention will now be described by way of example with reference to the accompanying drawings. In the drawings:

Figure 1 illustrates the operation of one example embodiment of a border detection algorithm on an image.

Figure 2 illustrates borders detected in the image of figure 1.

Figure 3 shows sets of adjacent pixels.

Figure 4 shows conditions in thresholded attributes of sets of adjacent pixels.

Figure 5 illustrates the tracing of a border on an image.

Figure 6 shows a border formed of interpolated or averaged points.

Figure 7 illustrates the operation of a border detection algorithm on another image.

Figure 8 shows a device suitable for image processing.

In one example algorithm, an image is composed of a set of pixels arranged on an orthogonal grid. A one-by-two block of adjacent pixel positions in the image is analysed to estimate whether a border passes through it. If no border is detected then the analysis steps in an arbitrary direction to a new one-by-two block until a border is found. When a border is detected in a one-by-two block, the analysis forms a two-by-two block including that one-by-two block. It estimates the direction of the border through the one-by-two block and based on the direction of the border steps to an adjacent two-by-two block in a direction along the estimated detected border. That block is analysed in turn, and in a similar way analysis steps along the border until the border has been fully traversed or the edge of the image has been reached. As will be described below, this algorithm has been found to be capable of being implemented in a highly efficient manner.

In one embodiment the present invention provides a method for reliably detecting borders/contours in an image without processing the entire image.

An image may be considered to be formed of a grid of pixels. For any horizontal and vertical position x, y a "border function" $b(x,y)$ can be defined which returns a real number. The function $b(x,y)$ can be defined differently depending on the use case, but preferably satisfies the following constraints:

- for any given x, y and image, it produces the same result each time it is calculated; and
- its output is at least roughly continuous (relevant if sub-pixel-accurate border detection is desired).

The function $b(x,y)$ can be defined so that two neighbouring pixel positions for which one has a negative and the other a non-negative $b(x,y)$ – or more generally for which a binary condition imposed on $b(x,y)$ yields a different value for each of the two neighbouring pixel positions – are interpreted as having a border between them.

It is advantageous for the implementation of $b(x,y)$ to be faster to compute once it has already been computed for a neighbouring pixel position.

By scanning for a border and then following it when found, it is possible to map borders without computing $b(x,y)$ for every pixel position in the image.

In a first routine of the method the image is scanned to detect a border. This can be done by taking an arbitrary set of adjacent pixel positions, computing their $b(x, y)$, and checking whether the binary condition in respect of each differs, in which case there's a border crossing in between. The pattern in which sets of adjacent pixel positions are sampled can vary depending on the use case. However, it is commonly advantageous to scan along a vertical or horizontal line. Such lines may be spaced at intervals in the image. In this arrangement not every set of two directly adjacent pixel positions is scanned, and some borders may be missed. However it has been found that with suitable choice of the interval between scan lines (e.g. 32 pixels or less) no major features of the image are normally missed, and there is a great performance advantage over computing $b(x, y)$ for every pixel position in the image.

Once a crossing has been detected, the method proceeds to a second routine of tracing the border. Given two adjacent pixel positions with a border between them, the border can begin to be traced. Two guidelines for this process can be predetermined arbitrarily:

1. Left-hand or right-hand "bias". If this bias is maintained borders can be circumnavigated in a deterministic fashion, the tracing routine returning to where it started.
2. The direction in which to first follow the border. Note that in some use cases it is desirable not to circumnavigate the border to its starting point. In those cases it may be desirable to trace in both directions from the point where the border is first detected.

Suppose that the binary condition applied to the attributes of pixel positions results in either a negative or a positive value being attributed to each pixel position that is considered, and the border is to be followed clockwise around non-negative areas, meaning we always have non-negative values on the right and negative on the left. Suppose the bias is selected to be right-hand-biased. Tracing the border then follows this pattern:

1. Given two adjacent pixel positions s_0 and s_1 , with a negative and a non-negative $b(x, y)$ respectively;

2. We take the neighbouring set of two adjacent pixel positions n_0 and n_1 , which themselves are adjacent to s_0 and s_1 respectively, in the chosen direction of tracing.
3. If both n_0 and n_1 have non-negative values, we move left (with respect to the current direction). If both have negative values, we move right. If n_0 has a negative and n_1 has a non-negative value we move forward. Lastly, if n_0 has a non-negative value and n_1 has a negative value, our next direction is determined by the "bias" mentioned above; In this case we move right due to our chosen bias.
4. Lastly our next set of adjacent pixel positions s_0 and s_1 are chosen as follows. If we moved forward: n_0 and n_1 . If we moved left: s_0 and n_0 . If we moved right: n_1 and s_1 .
5. Go to step 1 and repeat for the new s_0 and s_1 .

When s_0 and s_1 once again are the same two adjacent pixel positions we started with, we have looped around and come back to the beginning.

During this process, for each iteration of the above steps a border point is recorded. This yields a set of points in sequence along the border, in this case in the clockwise direction around non-negative areas.

How each border point is recorded can differ depending on whether sub-pixel accuracy is desired. In a non-sub-pixel-accurate variant each border position can be considered to be at the average of the four pixel positions after step 2. If the pixels are considered to be at integer positions, each element of any border position will be an integer plus 0.5. If the border function $b(x,y)$ is roughly continuous and we desire sub-pixel-accurate border detection, we can instead for each iteration record a border point as follows. For each set of adjacent pixel positions s_0 and s_1 , we linearly interpolate their positions to estimate the real point of the zero-crossing of $b(x,y)$. So for instance if s_0 's value is -1 and s_1 's is 2, the resulting border point will be s_0 's position times $2/3$ plus s_1 's position times $1/3$.

Figure 1 shows an example image to illustrate general principles of the border detection process described in more detail below. The image contains regions 1, 2 which are distinct from a background field 3. The regions 1, 2 are considered to be bounded by borders. It is desired to detect those borders, for example for the purpose of characterising the image. An arbitrary location 4 on the image is selected, together with an initial default scan direction as shown at 5. The location can conveniently be near a corner of the image but it need not be. The initial scan direction is conveniently parallel to axes on which pixels of the image lie, but it need not be. The process involves stepping through the image from the starting location 4 in the initial scan

direction 5. At each step the local region of the image is analysed to assess whether it contains a border. This analysis will be described in more detail below.

In the example of figure 1 the scanning initially follows path 6. Once the edge of the image is met, at 7, scanning continues in the same direction but from a point on opposite side of the image and offset by a predetermined distance 8 perpendicular to the initial scan direction 5. Scanning then continues at 9 until a border is met at 10. A tracing process is then initiated to circumnavigate the detected border, as illustrated at 11. As the border is followed, the locations of points on the border are stored for later use. In one example, once tracing returns to the point 10 where the border was first detected it resumes from a point 12 which is at the extension of the path 9 on which the border was met, but on the far side of the border from the location where it was met. Scanning then continues in a similar fashion along paths 13 and 14. When the same border is met again at 15 it can be ignored and scanning is resumed from point 16 which is at the extension of the path 14 on which the border was met, but on the far side of the border from the location where it was met. By continuing in a similar fashion the process can detect the borders of regions 1 and 2 and generate a set of stored border locations as illustrated in figure 2. In other examples, scanning could continue in other places and/or directions in the image.

The offset 8 can be set depending on any one or more of the level of accuracy required, the expected detail in the image, the amount of time available to process the image and the processing capabilities of the apparatus on which the algorithm is running. For example, to speed up the processing a greater value could be selected for offset 8.

The present algorithm has two distinct routines. In a first routine the image is scanned until a border is detected. This scanning conveniently steps one pixel position at a time along a straight line, but the scanning could be in other than straight lines or in steps of more than one pixel position. At each step in the scanning process the algorithm determines whether a border is present. If a border is present it determines a direction in which to follow that border. It then executes a second routine, which circuits the border as illustrated at 11 in figure 1. The mechanism for doing this will now be described.

In the first routine the assessment of the presence of a border is made on the consideration of one-by-two blocks of mutually adjacent pixels, each pixel being at a respective pixel position. For example, figure 3 shows a set of pixels aligned on orthogonal axes. In a first iteration of the first routine the algorithm may have reached pixel positions 20 and 21. It may perform an analysis of the image based at those pixel positions to assess whether a border is present there.

If no border is present then it may step to the next pair of pixel positions in the current scanning direction, i.e. to pixel positions 21 and 24. In a second iteration of the first routine it performs an analysis of the image based at those pixel positions to assess whether a border is present there. If a border is determined to be present at pixel positions 21, 24 then the direction of that border is estimated and the process initiates repetitions of the second routine in order to trace around the border.

In the second routine two-by-two blocks of pixels are considered. For a first iteration in the second routine two pixel positions that are adjacent to the two current pixel positions and offset from them in the estimated direction of the border are selected. Those may be pixel positions 23 and 25. The algorithm performs an analysis of those pixel positions to assess the next direction of the border. For a second iteration in the second routine two pixel positions adjacent to the four current pixel positions and offset from them in the estimated direction of the border are selected (those may be pixel positions 28 and 29). In the second iteration in the second routine the algorithm performs an analysis of those pixel positions to assess the next direction of the border. The algorithm proceeds in this fashion, tracing the border until it returns to the initial set of two adjacent pixel positions, in which case the border has been circumnavigated. Then the second routine can end and the first routine can resume. At each step in the second routine the algorithm steps by a distance of one pixel in a direction parallel to one of the orthogonal pixel axes. A consequence of this is that for any pixel position that needs to be analysed, an adjacent pixel position has already been analysed immediately or almost immediately prior (except for the very first pixel position in the scanning routine). This can provide significant efficiency advantages. In order to perform analysis on a pixel position, information must be gathered about its attributes. The gathering of that information may consume considerable processing power. In the described process, the function $b(x, y)$ can be defined and implemented such that computing $b(x, y)$ for a given pixel position is significantly faster if it's already been computed for an adjacent position.

The attribute of a pixel position that is used to help detect the presence of a border may depend on the purpose for which the image is being analysed. For example, if borders are being sought between light and dark parts of the image then the attribute could be based on a brightness value of the pixel. If borders are being sought between regions on the basis of their blueness then the attribute could be based on the blue intensity value of each pixel. In practice, the attribute is usefully a more complex measure such as the second-order rate of change (derivative) of the Gaussian, e.g. with respect to the gradient direction, or the Laplacian of Gaussian, or an

approximation thereof, based on an image property such as the luminance for pixels around a pixel position in question. This will be described in more detail below.

As indicated above, in each step a one-by-two or two-by-two block of pixel positions is under consideration. The relevant attribute for each pixel in the block is considered, and is assessed against a predetermined threshold. For example, if the attribute is brightness then the threshold might be 50% of the available brightness scale. For each of the pixels in the block a determination is made as to whether the attribute for that pixel is above or below the threshold. It is convenient to form a Boolean value from each attribute and to consider a border to be present between two pixel positions if the Boolean value differs between them.

Figure 4 illustrates representative possibilities for the result of this analysis. Figure 4 shows two-by-two blocks of mutually adjacent pixels in which pixels above the threshold are shown shaded and pixels below the threshold are shown unshaded. It is assumed in each case that the initial tracing direction is to the right and the two-by-two blocks have been approached from the left-hand side. The tracing algorithm is configured to traverse the border with pixels above the threshold on the left-hand side. Figure 4a shows the situation where both pixels in the newest one-by-two block (i.e. the pair of pixels on the right of two-by-two block 4a) are above the threshold. In that situation the tracing algorithm turns right, as shown at 40. Figure 4b shows the situation where one pixel in the newest one-by-one block is above the threshold and one below, and where each of those is adjacent to a pixel in the pair of pixels maintained from the previous iteration (i.e. the pair of pixels on the left of two-by-two block 4b) of the same type. In that situation the tracing algorithm continues in the same direction as before, as shown at 41. Figure 4c shows the situation where both pixels in the newest one-by-two block are below the threshold. In that situation the tracing algorithm turns left, as shown at 42. Figure 4d shows the situation where one pixel in the newest one-by-two block is above the threshold and one below, and where each of those is adjacent to a pixel in the pair of pixels maintained from the previous iteration of the opposite type. In that situation the tracing algorithm continues according to the predetermined bias, either turning left as shown at 43 or turning right as shown at 44. When the current one-by-two block yields a situation of the form shown in figure 4d, consisting of two non-adjacent pixels above the threshold and two non-adjacent pixels below the threshold, that might indicate a bifurcation of the border or a meeting of two borders. Collectively, this behaviour implements the logic described in numbered step 3 above. In each case the tracing path adopts a new path that keeps pixels of the same type (above/below the threshold) on the same side of the border.

At each step in the first routine, which scans the image to detect a border, the current one-by-two block is assessed. If no border has been detected in the current one-by-two block then the scanning steps one pixel in the current scan direction and the next one-by-two block is assessed, as described above with reference to figure 1. As indicated above, the new one-by-two block includes only one new pixel position so only one pixel position need be analysed.

In the first and second routines, when a border is detected then its location is estimated and stored. The location of the border can be estimated in a number of ways. The location that is stored may be the geometric centre of the one-by-two or two-by-two block under consideration, i.e. the mean of the pixel positions of that block: for example location 27 (see figure 3) for the two-by-two block consisting of pixels 21, 23, 24, 25. In another example the location of the border could be estimated to be the point where an interpolation between the determined attributes of the pixel positions of the current one-by-two block (or in other embodiments the current two-by-two block) would intersect a predetermined value. The interpolation may be a linear interpolation.

The direction in which the border is estimated to run is conveniently (i) parallel with an inferred border in the current two-by-two block and (ii) directed away from the centre of the current two-by-two block. Such directions are indicated at 40-44 in figure 4. One convenient approach is to treat the zones whose borders are to be traced as being regions above the threshold, the direction that is selected being that which has the pixel or pixels that are above the threshold on the left. In other embodiments the direction that is selected could be that which has the pixel or pixels that are above the threshold on the right.

The following of a border can stop when the tracing reaches a location where a border has already been detected. If that is due to the tracing reaching the point where the current border was first detected then the tracing of that border can be considered complete. The algorithm can then proceed to scan from any one-by-two blocks of the type shown in figure 4d that were detected during the traversal of that border, on the principle described above. Then the algorithm can continue from the far side of the border, as indicated at 12 in figure 1, in order to detect other borders. If the tracing of a border meets the edge of the image or another previously stored border then tracing can resume in the opposite direction from the point at which the currently traced border was first met. Tracing in the opposite direction can conveniently be achieved by adopting the opposite sense for the selection of a new tracing direction: for example

with pixels above the threshold on the left, as opposed to with pixels above the threshold on the right. This can allow the algorithm to fully trace borders that do not form complete loops.

An example of the border tracing process will be described with reference to figure 5. Figure 5 shows an orthogonal grid of pixels, in which pixels whose attribute is above the threshold are shown shaded. Pixels in figure 5 will be identified by a grid reference.

Scanning starts considering pixel positions A2 and B2, with the scanning direction being to the right as illustrated at 51. Pixel positions A2 and B2 are analysed. The attributes of both those pixel positions are on the same side of the threshold so the scanning direction is unchanged. Then pixel positions B2 and C2 are analysed. A border is detected between B2 and C2, directed up-and-down, so based on the predetermined choice of keeping attributes above the threshold on the right-hand side as we follow the border, next pixel positions B1 and C1 are considered. Their attributes are below the threshold, so the next direction is determined to be rightward. Then pixel positions D1 and D2 are considered. The algorithm proceeds in this way until the border has been circumnavigated.

At each step where a border is detected a border location is stored. In this example the border locations are stored as the geometric centres of each two-by-two block of pixel positions, the first two in the example above being $(B1 + C1 + B2 + C2) / 4$ and $(C1 + D1 + C2 + D2) / 4$. These locations are indicated by open circles as shown in figure 5.

As indicated above, the attributes of only two pixel positions need be computed at each step in the second routine. The attributes of the other two pixel positions of a two-by-two block can be stored from the previous step. This can help to make the algorithm highly efficient. The attributes of any pixel position previously considered can be stored indefinitely but it has been found that it is relatively efficient for computation and storage to retain attributes for only the two pixels that two successively-considered two-by-two blocks have in common.

It is convenient to store a position at each step around the border so that the border can be processed later and/or so that it can be known when the border has been fully mapped. One way to do this is to store at each such position the mean of the centres of the latest two pixels between which the border has been detected. In the examples of figure 4, this would be the lower two pixels in figure 4a, the right-hand two pixels in figure 4b and so on. Another option is to interpolate between the centres of the pixels based on an attribute of the pixels, most

conveniently the same attribute as the binary decision was applied to in order to determine the presence of the border. For example, suppose one pixel position is $x=20$, $y=40$ and has attribute -9 and the adjacent pixel position is at $x=21$, $y=40$ and has attribute 2. The binary test may be whether the attribute exceeds a threshold, such as zero. On that basis a border can have been determined between these pixel positions. The mean position would be 20.5,40. The interpolation may estimate where at a sub-pixel level a straight line having value -9 at position 20,40 and value 2 at position 21,40 would intersect a predetermined value. Conveniently that is the same predetermined value as is used in the binary test, in this case zero, which would mean the interpolated position would be 20.82,40. Figure 6 illustrates the effect of this in practice. Figure 6 shows a set of pixels including a region 71 of shaded pixel positions whose attributes are above a predetermined threshold. The path made up of line segments some of which are designated 73 shows the border of the region 70 determined as a route extending between interpolated points 72 on the centres of successive two-by-two blocks traversed around the border of the region. The path made of segments 73 shows the border of the region 71 determined as a route extending between points (some of which are indicated as 72) interpolated or averaged among the attributes of the pixel positions of each one-by-two block in which a border has been detected around region 71. As indicated above, it has been found that a path formed by joining points determined from the raw or non-thresholded attributes of the pixel positions, for example interpolated, is considerably more representative of the content of an image than a path formed from raw pixel positions.

It is notable that the process of forming a border from interpolated or averaged points involves detecting the existence of a border and following it based simply on whether the relevant attribute of the pixel positions under consideration exceed a threshold, but estimating the location of the border based on the actual value of the attribute for each of those pixels. This approach has been found to be relatively computationally efficient and also relatively representative of the image content.

The attribute of each pixel position that is used in the processing described above can be selected dependent on the characteristics it is desired to analyse. In some examples, the attribute could be dependent solely on the values of that individual pixel: for example its brightness or hue. In other examples the attribute of the pixel can be determined by a function that is also dependent on the values of neighbouring pixels. In this latter case the algorithm may be variable so as to allow the relative influence of pixels distant from the pixel in question to be increased or decreased. The attribute may be determined by summing intermediate values calculated for a set

of pixels neighbouring the pixels in question. The extent of those neighbouring values may be variable. The extent may be selected in dependence on the level of detail, focus or range at which the image is being considered. Each intermediate value may be weighted by a weighting factor determined according to a weighting algorithm. In one preferred example the attribute of each pixel is determined by means of the Laplacian of Gaussian filter. The Laplacian of Gaussian can be calculated by determining a weighted sum of the intensities of pixels neighbouring the pixel in question, the weights being determined according to the Laplacian of Gaussian filter centred on the pixel in question. With the pixel in question at $x=0$, $y=0$, the Laplacian of Gaussian filter gives a weighting for a pixel at x , y of:

$$-\frac{1}{\pi\sigma^4}\left(1 - \frac{x^2 + y^2}{2\sigma^2}\right)e^{-\frac{x^2+y^2}{2\sigma^2}}$$

Where σ represents the Gaussian standard deviation, influencing the weighting of the function with radius from the pixel in question. The weightings may be calculated in advance and stored in a look-up table for use when the attribute of a given pixel is to be computed. Weightings for a range of values of σ may be pre-computed and stored and applied as appropriate when the image is to be analysed at an appropriate level of detail. The weightings may be approximated for convenience. When estimating the attribute of a particular pixel position the weightings may conveniently be applied only to the pixels of a square block (i.e. a block whose sides are of equal pixel length) centred on that pixel position, and thus not to pixels outside that block. This can simplify the computation.

The Laplacian of Gaussian is a convenient filter for use in the present system because it tends to emphasise borders. Other filters with similar properties, including but not limited to difference-of-Gaussians, determinant of Hessian, Sobel, Canny and Canny-Deriche could be used either alone or in combination with other functions to determine the attribute of a pixel position.

Figure 7 shows another example image to illustrate the border detection process according to another embodiment. In this embodiment, the determination of whether or not to map a border location at each tracing step is dependent on the strength or quality of the border at that step. Mapping, for example, could involve using that location to identify features in an image and/or storing of the border location.

At each step in the scanning process the algorithm determines whether a border is present. The assessment of the presence or direction of a border can also include a determination of the strength or quality of the detected border. Based on this determination, a border may, for example, be classified as high or low quality and the algorithm may decide that only high quality borders be considered. This can help improve feature detection as features such as shadows and noise may have poor quality borders and thus can be discarded.

Similarly to figure 1, the image of figure 7 contains regions 1, 2 which are distinct from a background field 3. In addition, figure 7 also contains a line 90. Line 90 may, for example, be an object in the image or an artefact or noise created during image capture or processing. Line 90 differs from regions 1 and 2 in that the border for line 90 does not form a continuous perimeter around a region.

As described above, the image is scanned until a border is met at 10. The border is then traced, as illustrated at 11. At each tracing step, the quality of the border between regions 1 and 3 is assessed. In this example, the quality of the border is high and thus the border is mapped and followed. Once tracing returns to the point 10 where the border was first detected, it is determined that the mapped path forms a perimeter around region 1 and scanning resumes from a point 10 which is the point at which the border was first detected.

A border is met again at 91 and the border is traced and followed in a similar direction, as illustrated at 92 as the previous border at 10. The tracing reaches point 93, where it is determined that the border quality is poor. The border from this point is no longer traced and is not mapped. The tracing then returns to point 91 and follows and maps the border in the opposite direction, as illustrated at 94, to 95, where line 90 ends. The scanning for further borders then resumes from point 91 in the scanning direction.

The quality of the border determined at each step can be dependent on the attribute of the pixel position or positions that is used to help detect the presence of the border. For example, if the attribute is the brightness value of a pixel position, the amount of change in brightness of pixel positions that are perpendicular to the border could be used to determine the border quality. A first-order derivative could be estimated from those pixels to calculate a gradient magnitude, which can then indicate the quality of the border. The gradient magnitude could be assessed against a predetermined threshold to determine if the border is of sufficient quality. The quality determination can be applied to other attributes, such as those described herein.

As described above, the assessment of the presence or direction of a border can be made on the consideration of a block of two mutually adjacent pixel positions. If one of the two adjacent pixel positions has an attribute value that is above a predetermined threshold and the other has an attribute value that is below the predetermined threshold, this can indicate that there is a border present somewhere between the two pixel positions. For example, the attribute of the pixel position could be based on an estimate of the second-order derivative of an isotropic Gaussian convolution of the image with respect to the gradient direction of the Gaussian. The border quality measure can be based on a third-order derivative of an isotropic Gaussian convolution of the image with respect to the gradient direction of the Gaussian for each of the two pixel positions. In this case, the third-order derivative is a scalar value where a higher value indicates a better-quality border.

To simplify the border detection and storage the system may ignore borders the size of which are below a pre-determined threshold. That size may be defined as the border circumference or length. If a border once circumnavigated proves to be below the pre-defined minimum size then it is not stored and is discarded.

As each border is mapped, points along its length are stored. This allows it to be determined later on whether a border has already been found, and once the algorithm has completed it provides a record of the borders in the image.

Once the borders have been located any suitable processing can then be undertaken.

In a first example, the image can be modified so as to highlight the locations of the detected borders, e.g. in an eye-catching colour. Then the resulting image can be displayed. This may be useful in a vehicle for highlighting the borders of a road viewed through a camera mounted on the vehicle.

In a second example, the regions of the image that correspond to an area can be transformed, and the resulting image can be displayed. The transformation is applied to the region within the outer border of the area but outside any inner borders of the area. The transformation may, for example, be to lighten or darken that area, to apply a screen (e.g. cross-hatching) to that region, or to change the chrominance of that region. If the region is lightened or darkened then the chrominance at each point in the region may remain unchanged. This can give the impression of shining a light on, or shading, the region. If the chrominance of the region is changed then the

luminance at each point in the region may remain unchanged. This can give the impression of re-colouring the region. This is useful, for instance, for simulating painting the region in a certain colour.

Using the methods described herein each frame of the video can be processed, transformed and displayed in real time.

Figure 8 illustrates apparatus for implementing the methods described above. The apparatus comprises a housing 80 in which are a camera 81, a display 82, a processor 83, a non-volatile memory or ROM 84 and a working memory or RAM 85. The apparatus is powered by a battery 86. In this example the display 82 is a touchscreen, so it provides user input to the processor 83 as well as being driven by the processor, but a separate keypad could be provided. The ROM 84 stores program code that is executable by the processor. The program code is stored in a non-transient form. The program code is executable by the processor to perform the functions described above. In operation the processor can receive an image, either from the camera 81 or from a remote communications transceiver 87. In the former case, the image could be an image captured by the camera of the environment at the location of the device. In the latter case, the image could be downloaded from the internet. The processor stores the image in RAM 85. The image could be displayed on the display 82. Once the image is stored in RAM, the processor can analyse it using an algorithm as described above. Then, dependent on the outcome of the analysis the process may perform further operations in dependence on the stored program code. For example, it could alter the image to highlight detected borders. Alternatively it could re-colour regions inside one or more detected borders, for example in response to selection of those regions by a user by means of the touchscreen. Alternatively, a user could select one of those regions, and the processor could identify other regions that are surrounded by detected borders and have attributes in common with the selected region (for example a chrominance that differs by less than a predetermined threshold from the average chrominance of the selected region) and then recolour both the selected and the identified regions similarly. Alternatively, the processor could use the position of the detected borders to infer information about the three-dimensional spatial relationship of objects represented in the image.

In the example above, the device can capture image data and process it locally. Alternatively, it could capture image data, transmit it to a remote server for analysis and then receive back from the server information related to the detected borders.

The device of figure 8 could be a mobile phone.

The image could be a still image or a frame or portion of a frame of a video stream.

Each pixel is a sub-unit of the image. The pixels could be monochrome or colour pixels. For display or image capture a colour image is often divided into channels of individual colours which when considered collectively represent a certain luminance and chrominance at a certain location. For the purpose of performing the algorithm described above a single pixel of the image may be considered to be formed by those channels in combination. Preferably for the purpose of performing the algorithm described above each pixel fully characterises the visual content of the image at the location of the pixel.

The applicant hereby discloses in isolation each individual feature described herein and any combination of two or more such features, to the extent that such features or combinations are capable of being carried out based on the present specification as a whole in the light of the common general knowledge of a person skilled in the art, irrespective of whether such features or combinations of features solve any problems disclosed herein, and without limitation to the scope of the claims. The applicant indicates that aspects of the present invention may consist of any such individual feature or combination of features. In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention.

CLAIMS

1. A method for identifying features in an image, the method comprising:
 - a) from a first location, scanning the image along a scan path in a scan direction away from the first location so as to detect a border in the image at a second location;
 - b) mapping a path defined by the border in a first direction along the border from the second location, the path being mapped in dependence on a border quality measure;
 - c) determining whether the mapped path forms a perimeter; and
 - d) if the mapped path does form a perimeter, scanning the image from the mapped path so as to detect a further border in the image; or if the mapped path does not form a perimeter, mapping a path defined by the border from the second location in a second direction along the border that is different to the first direction.
2. A method as claimed in claim 1 comprising:
 - on mapping said path in the second direction, scanning the image from the mapped path so as to detect a further border in the image.
3. A method as claimed in claim 1 or 2, wherein the method comprises, on detecting a further border in the image at a third location, repeating steps b) to d), wherein the repeated steps are performed from the third location.
4. A method as claimed in any preceding claim, wherein the method comprises, after having traced a border:
 - determining whether the circumference or length of the border and/or the compounded quality measurements of the border exceeds a predetermined threshold; and if not, discarding the border.
5. A method as claimed in any preceding claim, wherein the method comprises, from a fourth location: scanning the image along a second scan path in the scan direction away from the fourth location so as to detect a border in the image at a fifth location, the fourth location being offset by a predetermined distance from the first location in a direction perpendicular to the scan direction.
6. A method as claimed in any preceding claim, the method comprising, on detecting the border:
 - ignoring the border if it has been mapped previously.

7. A method as claimed in any preceding claim, wherein in each step of scanning the image, the scan direction is the same.

8. A method as claimed in any preceding claim, wherein the image is composed of a plurality of pixels and the method comprises mapping a path by detecting borders between pixel positions, wherein the quality measure is indicative of the quality of the border at each border position.

9. A method as claimed in claim 8, wherein the method comprises:

estimating an attribute for each of a set of two adjacent pixel positions in the image;

assessing whether a predetermined binary condition differs in respect of the two pixel positions, and if so, determining that a border is present in the part of the image represented by the two pixel positions.

10. A method as claimed in claim 9, wherein the border quality measure is a value derived from neighbouring pixels, the method further comprises: determining whether the derived value is above or below a predetermined threshold, and if above, mapping that border position as part of the path; or if below, not mapping that border position as part of the path.

11. A method as claimed in claim 10, wherein the attribute of the pixel position is based on a Laplacian of Gaussian filter applied to the image and centred on that pixel position.

12. A method as claimed in claim 9, wherein the attribute of a pixel position is based on estimating the second-order derivative of an isotropic Gaussian convolution of the image at the pixel position and with respect to the gradient direction of the Gaussian.

13. A method as claimed in claim 12, wherein a second attribute of a pixel position is based on estimating the third-order derivative of an isotropic Gaussian convolution of the image at the pixel position and with respect to the gradient direction of the Gaussian, wherein the border quality measure at a given border position is an interpolation of the second attribute with respect to each of the two pixel positions on either side of the border position.

14. A method as claimed in any preceding claim, wherein the image is a frame of a video stream.

15. A device comprising:

a memory storing non-transient program code; and

a processor configured to identify features in an image by executing the code to perform the steps of:

a) from a first location, scanning the image along a scan path in a scan direction away from the first location so as to detect a border in the image at a second location;

b) mapping a path defined by the border in a first direction along the border from the second location, the path being mapped in dependence on a border quality measure;

c) determining whether the mapped path forms a perimeter; and

d) if the mapped path does form a perimeter, scanning the image from the mapped path so as to detect a further border in the image; or if the mapped path does not form a perimeter, mapping a path defined by the border from the second location in a second direction along the border that is different to the first direction.

16. A device as claimed in claim 15, wherein the device comprises a camera and the image is an image captured by the camera.

17. A device as claimed in claim 16, wherein the image is a frame of a video stream captured by the camera.

18. A device as claimed in claim 17, wherein the device is configured to perform the said steps in real time on successive frames of the video stream.




FIG. 1




FIG. 2

2 / 4

		27	
<u>20</u>	<u>21</u>	<u>24</u>	
<u>22</u>	<u>23</u>	<u>25</u>	
	<u>28</u>	<u>29</u>	

FIG. 3




FIG. 4a




FIG. 4b




FIG. 4c






FIG. 4d

3 / 4

*FIG. 5**FIG. 6*

4 / 4




FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2015/051405

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T7/00
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, COMPENDEX, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y A	US 4 910 786 A (EICHEL PAUL H [US]) 20 March 1990 (1990-03-20) abstract column 6, lines 38-53 column 10, lines 15-22 column 11, lines 36-41 column 12, line 62 - column 13, line 14 column 17, lines 6-43 column 18, lines 22-41 column 22, lines 49-51 column 23, line 11 - column 24, line 11 claims 1-3 ----- -/--	1-5,7,8, 14-18 6,9-13



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

29 September 2015

Date of mailing of the international search report

07/10/2015

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Eveno, Nicolas

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2015/051405

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	CHANG F ET AL: "A linear-time component-labeling algorithm using contour tracing technique", COMPUTER VISION AND IMAGE UNDERSTANDING, ACADEMIC PRESS, US, vol. 93, no. 2, 1 February 2004 (2004-02-01), pages 206-220, XP004483599, ISSN: 1077-3142, DOI: 10.1016/J.CVIU.2003.09.002	1-5,7,8, 14-18
A	page 207, paragraph 3 page 209	6,9-13
Y	----- WO 2009/013636 A2 (UNIV WESTERN ONTARIO [CA]; SAIKALY MANALE [CA]; FENSTER AARON [CA]; HA) 29 January 2009 (2009-01-29) claim 1	4
A	----- Peter Selinger: "Potrace: a polygon-based tracing algorithm", 20 September 2003 (2003-09-20), pages 1-16, XP055149450, Retrieved from the Internet: URL: http://potrace.sourceforge.net/potrace.pdf [retrieved on 2014-10-28] section 2.1	1-18
A	----- FARAG A A ET AL: "Edge linking by sequential search", PATTERN RECOGNITION, ELSEVIER, GB, vol. 28, no. 5, 1 May 1995 (1995-05-01), pages 611-633, XP004011490, ISSN: 0031-3203, DOI: 10.1016/0031-3203(94)00131-5 sections 2, 4.2, 5	1-18
A	----- CLARK J J: "AUTHENTICATING EDGES PRODUCED BY ZERO-CROSSING ALGORITHMS", IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE COMPUTER SOCIETY, USA, vol. 11, no. 1, 1 January 1989 (1989-01-01), pages 43-57, XP000032906, ISSN: 0162-8828, DOI: 10.1109/34.23112 section II -----	1-18

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2015/051405

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 4910786	A	20-03-1990	NONE

WO 2009013636	A2	29-01-2009	US 2010134517 A1 03-06-2010
			WO 2009013636 A2 29-01-2009
