



(51) International Patent Classification:

G06K 9/00 (2006.01)

(21) International Application Number:

PCT/EP2017/075010

(22) International Filing Date:

02 October 2017 (02.10.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: HP INDIGO B.V. [NL/NL]; Startbaan 16, 1187 XR Amstelveen (NL).

(72) Inventors: CHEN, Eli; Einstein 10 Kiryat Weizmann, 76101 Ness Ziona (IL). HAIK, Oren; Einstein 10 Kiryat Weizmann, 76101 Ness Ziona (IL). PERRY, Oded; Einstein 10 Kiryat Weizmann, 76101 Ness Ziona (IL).

(74) Agent: HGF LIMITED; Fountain Precinct Balm Green, Sheffield South Yorkshire S1 2JA (GB).

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))

(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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,

(54) Title: VALIDATION OF CANDIDATE LOCATIONS OF A BOUNDARY



(57) Abstract: A method is described in which an image is received, the image including a document and background area; a candidate location of a boundary between the document and the background area is found; and a validity of the candidate location of the boundary is checked, based on pixel values around the candidate location of a first channel of the image, and pixel values around the candidate location of at least one other channel of the image, wherein the at least one other channel is different from the first channel.

Figure 3

VALIDATION OF CANDIDATE LOCATIONS OF A BOUNDARY

BACKGROUND

[0001] Corner detection may be used to locate corners of a document that is contained in an image. Knowing the location of the corners allows the document to be registered, and the document extracted from the image. Extraction of a document from an image may find application in digital printing technologies, as well as other fields, such as image processing.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Examples are further described hereinafter with reference to the accompanying drawings, in which:

Figure 1 illustrates an image in accordance with some examples.

Figure 2a illustrates a method in accordance with some examples.

Figure 2b illustrates a method to find the boundary of a document in accordance with some examples.

Figure 3 illustrates a method to check a validity of a candidate location of a boundary in accordance with some examples.

Figure 4 illustrates a method to check a validity of a candidate location of a boundary in accordance with some examples.

Figure 5 illustrates a method to identify of a corner of a document in accordance with some examples.

Figure 6 is a schematic of the method to improve the identification of a corner of a document in accordance with some examples.

Figure 7a illustrates an apparatus to find a validity of a candidate location of a boundary in accordance with some examples.

Figure 7b illustrates an apparatus to refine a corner location estimate in accordance with some examples.

Figure 8 illustrates a system to find a validity of a candidate location of a boundary in accordance with some examples.

Figure 9a illustrates a computer readable medium in accordance with some examples.

Figure 9b illustrates a computer readable medium in accordance with some examples.

Figure 10 illustrates a printing device 1000 according to some examples.

DETAILED DESCRIPTION

[0003] Figure 1 illustrates an image 100. The image 100 may be captured by an image

capture device (e.g. produced by an image scanner) and may be referred to herein as a captured image 100. In some examples the captured image 100 is produced by an inline scanner (e.g. of a printing apparatus, such as a digital printing press), however the image scanner may be any type of image scanning apparatus or other image capture apparatus, such as a digital camera. The captured image 100 may comprise a document 120 and a background area 130. When capturing an image of a document 120 the image scanning apparatus may also capture the surrounding area of the scanner backplate or surface the document 120 is placed on or against. This background surface may form the background area of the captured image. A user may wish to extract the document 120 from the captured image, or otherwise identify the document. The document 120 may correspond with the document page or similar that a user desires to extract from the captured image 100. The document 120 may be skewed relative to the captured image 100. In some examples, for example when using an inline scanner, a smaller skew may be expected than that shown in figure 1, which has been selected for illustrative purposes. In some examples, the document may be square or rectangular.

[0004] In some cases, detection of the document corners 120a-120d may be challenging if the contrast between the document corners 120a-120d and the background area 130 is low. This may be the case when the document 120 comprises a non-white substrate, or when the background area 130 and the document 120 have a similar color. Some examples may provide improved corner detection in such cases.

[0005] Figure 1 also illustrates an image processing apparatus 700 to receive the captured image 100, and process (e.g. extract) the document 120 according to the examples herein.

[0006] According to some examples, the document 120 may be extracted by identifying locations of the document corners 120a-120d. A registration process may be performed on the captured image 100, such that the captured image 100 is translated and/or rotated into a coordinate system in which the document 120 is aligned with reference axes, based on the location of the document corners 120a-120d. In some examples the reference axes may be aligned with features of the image scanner, or the edges of the captured image prior to the registration. For example, following registration, one document corner 120c may be mapped to the origin of the reference axis (0,0), document corner 120b may be mapped to point (b,0) relative to the reference axis, document corner 120a may be mapped to point (b, a) relative to the reference axes, and document corner 120d may be mapped to point (0,a) relative to the reference axes.

[0007] Once the document 120 is registered it may be extracted from the page, in some examples this may be performed by cropping the captured image to obtain the document 120.

Improving the accuracy of the location of the document corners 120a-120d may lead to an improvement in the registering of the document 120 and/or cropping of the captured image 100. The following examples describe techniques which may be used to find the location of the document corners 120a-120d. In some examples the document may be cropped prior to registration. In some examples, cropping and registering the document may be performed simultaneously (e.g. in a single transformation).

[0008] The captured image 100 may be represented as a plurality of pixels. The pixels may be arranged in a regular two-dimensional grid. Each pixel may have one or more corresponding values describing of the appearance of that pixel. For example, each of the one or more values of the pixel may indicate an intensity of a particular color, a luminance, a chrominance, lightness, etc. The captured image may be described by one or more channels. Each channel corresponds with a value associated with a pixel to describe the appearance of the pixel. For example, if the captured image is represented in a RGB color space, the image has a red channel, a green channel and a blue channel, and each pixel has a respective red value, green value and blue value representative of the intensity of the respective color (red, green or blue) in the pixel. Other color spaces could be used, such as CMYK (cyan, magenta, yellow and black channels), HSV (hue, saturation and value channels), HSL (hue, saturation and lightness/luminance channels), LAB (lightness, green-red and blue-yellow channels) etc.

[0009] A grayscale image is an image with a single channel; that is, an image in which each pixel has a single respective associated value. Each channel of an image having multiple channels can be viewed as a respective grayscale image. Herein, grayscaling refers to generating a grayscale image from an image having one or more channels. The result of grayscaling an input image is referred to herein as a grayscaled image. The grayscaled image may be considered to be an additional channel of the input image.

[0010] In some examples, the respective value of each pixel of a grayscaled image may correspond with a luminance of the corresponding pixel of the input image. The luminance may be, for example in an RGB color space, a weighted sum of the intensities of each color of a pixel. In some examples, the luminance of the captured image 100 may be defined as

$$\text{[0011]} \quad L = a * R + b * G + c * B \quad (1)$$

where L is the luminance, R, G, and B are the respective intensities of each channel of the color image, and a, b and c are weighting constants. The parameters a, b and c may be chosen such that L corresponds with physiological perceived colors in human color vision, for example a = 0.30, b = 0.59, and c = 0.11 may be used. Luminance may be derived from other color spaces, either via a direct transformation or by converting to one or more intermediate color spaces.

[0012] The captured image 100 may be a color image, black and white image, or grayscale image. In some examples, each channel (e.g. a channel of an input image or a grayscaled image) may have 4 bits of information per pixel (e.g. 16 levels of intensity). In some examples, each channel may have 8 bits of information per pixel (e.g. 256 levels of intensity). In some examples, each channel may have 16 bits of information per pixel (e.g. 65,536 levels of intensity). Other numbers of bits of information per pixel may be used.

[0013] Figure 2a illustrates a method 200 to identify the document corners 120a-120d of a document 120. The method 200 may be performed by image processing apparatus 700. A captured image 100 may be converted into a grayscaled image at 210. At least one candidate location of a boundary between the document and the background area may be found at 220. The boundary may be found based on the grayscaled image, which in some examples may reduce processing involved in finding the at least one candidate location. The validity of the at least one candidate location may be checked at 230 by a validation process. If the candidate location is determined to be valid, then the candidate location may be used to define a bounding box at 240, otherwise a new candidate location may be found. The bounding box provides an estimate of the document 120 location, and gives an approximate (or estimated) location for each of the document corners 120a-120d of the document 120. A window may be defined 250 at or around each corner of the bounding box 110a-110d. The operations at 220, 230, and 240 may be repeated at 260 for each of the windows to provide a refined window for each of the document corners 120a-120d, thus providing a refined approximation (or estimate) of the location of each of the document corners. An example method of refinement of the windows is described in relation to figures 7 and 8.

[0014] A corner detection method may be applied within each of the refined windows at 270 to identify one or more potential corners in each window. In some examples, applying the corner detection within the windows may lead to a reduction in processing burden, compared with applying corner detection over the whole image. In some examples, applying the corner detection within the windows may lead to improved accuracy (e.g. fewer false positives) relative to applying corner detection across the whole image.

[0015] A selection may be made from among the identified potential corners at 280 to identify a reduced subset of the potential corners. A respective corner may be selected from each reduced subset, and designated a detected corner of the document. Accordingly, the location of each of the document corners 120a-120d may be approximated as a respective the location of the corresponding detected corner (being one of the potential corners of the corresponding reduced subset).

[0016] In some examples, the potential corner of the reduced subset having the shortest

distance relative to the corresponding corner of the bounding box is selected as the detected corner. For example, the top left detected corner is the potential corner (of the reduced subset of the top left window) that is closest to the top left corner of the bounding box.

[0017] The scanned image may then be processed 295 in accordance with the detected document corners. For example, by showing the captured image 100 on a display with the location of the detected document corners marked or identified. In some examples, the processing 295 includes cropping and/or registering the document 120.

[0018] The conversion of the captured image 100 at 210 may be based on the conversion factor in equation (1). Alternatively, it may be based on another conversion factor. The captured image 100 may be converted into grayscale before or after being received by the image processing apparatus 700. In some examples the conversion may be based on a single channel of a color image, for example a green channel of an RGB image may be used. The human eye is more sensitive to green than red or blue, and so the green channel gives an approximation of luminance without requiring the processing to calculate the luminance according to equation (1).

[0019] Finding at least one candidate location of a boundary at 220 may be based on a standard deviation of a property of the pixels of each row and column of the captured image 100. In some examples, the property may be a luminance of, or derived from, the captured image (e.g. by grayscaling the image). This is described in more detail with reference to figure 2b. In some examples, the at least one candidate location is found based on an identification of first 225 and last 226 discontinuities in the standard deviation of the property of the rows of pixels 221 of the captured image 100. The candidate location may also be based on the first 227 and last 228 discontinuity in the standard deviation of the property of the columns of pixels 223 of the captured image 100. In some examples, discontinuities in the standard deviation of the property may be identified using an edge detection technique, for example a Canny edge detector. A discontinuity may be a sharp change in value of the standard deviation. The property may be a gray level, a luminance, or an intensity (such as an intensity of the green channel of an RGB image) although other image properties that exhibit a discontinuity at a document edge could be used.

[0020] In some examples, a standard deviation of the intensity of the rows of pixels may be calculated, such that each row has a corresponding standard deviation, being the standard deviation of the property values of the pixels in that row. The first and last discontinuities in the standard deviation of the rows of pixels may be identified as candidate locations of the boundary. Here, first and last may be with respect to a direction perpendicular to the rows. For example, if the rows are numbered consecutively, the first discontinuity may correspond with

the lowest numbered row that corresponds with a discontinuity in the standard deviation.

Similarly, the last discontinuity may correspond with the highest numbered row that corresponds with a discontinuity. Discontinuities may be detected using a Canny edge detector. For example, the Canny edge detector may be applied to a graph of standard

5 deviation of the property (SD) against row number. In some examples, location in a column direction may be used in place of row number. The first and last discontinuities may determine position, along the column direction, of the candidate locations.

[0021] In some examples, a standard deviation of the intensity of the columns of pixels may be calculated, and the first and last discontinuities in the standard deviation of the columns of
10 pixels may be identified using a Canny edge detector. The first and last discontinuities of both the rows and columns may identify the candidate locations. The candidate locations may define the boundaries at or between intersections between the rows and columns having the first and last discontinuities. That is, if rows a and b have the first and last discontinuities, respectively, and columns c and d have the first and last discontinuities, respectively, the
15 edges between (c,a) and (d,a), (d,a) and (d,b), (d,b) and (c,b), and (c,b) and (c,a) may be identified as candidate boundaries.

[0022] In some examples the validity of the candidate location may be checked at 230. In some examples this may be based on pixel values a first channel of the captured image 100
20 around the candidate location, and pixel values of at least one other channel of the captured image 100 around the candidate location, wherein the first channel is different from the at least one other channel.

[0023] An example of the method to check the validity at 230 is described in in more detail below in relation to figure 3.

[0024] The bounding box 110 may be defined at 240 based on the candidate locations. The
25 bounding box corners 110a-110d of the bounding box 110 may not be in the same location as the document corners 120a-120d, for example if the document 110 is skewed relative to the background area 130.

[0025] A window may be defined around each boundary location or corner. In some examples, the size of the window may be a percentage of the size of the captured image 100.
30 In some examples, the window may be a constant size (e.g. a predetermined number of pixels). In some examples the window may have the same aspect ratio as the captured image 100, or in some examples the window may have a different aspect ratio. The window may be square, or rectangular. In some examples the window may be circular or oval. In some examples the window may have an area of 10 pixels by 10 pixels. In some examples the
35 window may have an area of $x \times y$ pixels where x is less than the width of the captured image

100 in pixels and y is less than the height of the captured image 100 in pixels. In some examples the window may contain or be centered on the candidate location. In some examples a corner of the window may be aligned with a boundary of the document 120, such that the corners of the boundary and window coincide. In some examples, respective windows
 5 may be defined for each candidate location. In some examples a corner of each window may be aligned to a corner of the bounding box 110 such that each window is substantially inside the bounding box 110.

[0026] The bounding box corners 110a-110d may be refined at 260 by a refinement process. An example of this process is described in more detail below, with reference to figures 7 and 8.

10 The refinement process may comprise finding an updated candidate location, validating the updated candidate location, and updating the window based on the updated candidate location. This process of refinement may be performed a set number of times, or alternatively may be performed until a predetermined criterion is met. In some examples the criterion may include the updated candidate location and previous candidate location being within a certain
 15 distance (e.g. a user-specified limit) of each other.

[0027] Corner detection may be applied within each window at 270 to identify a plurality of potential document corners in each window. In some examples the corner detection may comprise applying a Harris corner detector. Other corner detection methods may be used, such as a SUSAN corner detector or a minimum eigenvalue detector. Corner detection may
 20 result in a plurality of potential document corners identified inside the window. A document corner 120a-120d may be identified based on a comparison of the candidate location and the positions of the plurality of potential corners. Each potential corner may have an associated strength assigned by the corner detection method. For example, a Harris corner detector associates a value of R with each potential corner that it detects, as described below.

25 **[0028]** A Harris corner detector assumes a change in intensity $E(u,v)$ for a shift $[u,v]$ is defined by

$$E(u, v) = \sum_{x,y} w(x, y) [I(x + u, y + v) - I(x, y)]^2 \quad (2)$$

[0029] Where $w(x,y)$ is a window function, $I(x+u, y+v)$ is the shifted intensity, and $I(x,y)$ is the non-shifted intensity. The window defined by the window function for the purposes of the

30 Harris corner detection does not correspond with the window used in operation 250 or 260. Herein, the window defined by the window function will be referred to as a Harris window, to distinguish these different windows. The window function may be defined as 1 inside the Harris window, and 0 outside, or alternatively may be a Gaussian window. Other window functions may also be used. For Harris windows that contain no boundaries, $I(x+u, y+v) - I(x,y)$ will be

small. However, if there is a boundary in the Harris window then $I(x+u, y+v) - I(x, y)$ will be large. Therefore a Taylor expansion may be used, to obtain

$$\sum [I(x+u, y+v) - I(x, y)]^2 \approx \sum [u^2 I_x^2 + 2uv I_x I_y + v^2 I_y^2] \quad (3)$$

[0030] where I_x and I_y are the partial derivatives of I in the x and y directions. Therefore $E(u, v)$ may be rewritten in matrix form as

$$\left(\sum \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \right) \begin{bmatrix} u \\ v \end{bmatrix} \quad (4)$$

[0031] And for small shifts $[u, v]$ a bilinear approximation may be used to define

$$E(u, v) \cong [u, v] M \begin{bmatrix} u \\ v \end{bmatrix} \quad (5)$$

[0032] Where

$$M = \sum w(x, y) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \quad (6)$$

[0033] I_x and I_y may then be plotted against each other for each point inside the Harris window to obtain a scatter plot. An ellipse covering the points of radii λ_1 and λ_2 may then be calculated. If λ_1 and λ_2 are small and approximately equal then there is likely no corner or boundary inside the Harris window. If λ_1 and λ_2 are large and approximately equal then there is likely a corner inside the Harris window. If λ_1 is large and λ_2 is small (or vice versa), then there is likely a boundary inside the Harris window.

[0034] The strength of the corner may be defined as R , where $R = \det M - k(\text{trace } M)^2$ where k is an empirically determined constant in the range of 0.04 to 0.06. It is noted that $\det M = \lambda_1 \lambda_2$ and $\text{trace } M = \lambda_1 + \lambda_2$.

[0035] The top-rated corners (e.g. the corners identified as strongest, having the highest values of R) may be selected at 280 from the potential corners to obtain a subset of potential corner locations. In some examples this may comprise selecting n potential corner locations with the highest strength, where n is any integer. In some examples n may be 5, 10 or 20.

[0036] The detected corner locations may be set at 290. In some examples, the corner locations may be identified as the potential corner location, among the subset of potential corner locations, inside the window that is closest to the corresponding bounding box 110 corner. For example, the upper right document corner 120a may be identified based on a

potential corner location (of the subset of potential corner locations) that is closest to the upper right bounding box corner 110a; the lower right bounding box corner 110b may be identified based on the potential corner location that is closest to the lower right bounding box corner 110b; the lower left bounding box corner 110c may be identified based on the potential corner location that is closest to the lower left bounding box corner 110c; the upper left document corner 120d may be identified based on a potential corner location that is closest to the upper left bounding box corner 110d.

[0037] In some examples the corner location may be set as the potential corner location, among the subset of potential corner locations, that is closest to the corresponding bounding box corner 110.

[0038] In some examples a corner location may be selected from potential corner locations by first selecting a subset that are closest to the corresponding corner of the window, and then selecting a corner among the subset that has a highest strength value. In some examples a metric may be determined combining, for example, distance from the corresponding corner of the window and a strength value. The metric may be a linear combination or in some other combination of factors to be used in selecting a single corner from the plurality of potential corner locations.

[0039] In some examples the window corners may have a different location from the document corners 120a-120d. For example, due to a skew of the document.

[0040] At 295, processing may be carried out on the image, based on the locations of the detected corners. For example, the location of the document corners 120a-120d may be marked or indicated on a display of the captured image. In some examples the location of the document corners may be stored. In some examples, processing 295 may include registering and/or cropping the captured image.

[0041] An example method 300 for validating a candidate location of a boundary between a document 120 and the background area 130 is described in accordance with figure 3. The method 300 may be used to validate a location of a boundary. The method 300 may be used in operation 230 of figure 2a. In some examples, the method 300 may be used in other applications or used independently.

[0042] The captured image 100 may be received at 310. The captured image 100 may include a document 120 and a background 130. A candidate location of a boundary between the document 120 and the background 130 may be found or received at 320. The validity of the candidate location may be determined at 330.

[0043] In some examples, at least one candidate location may be found at 320. In some

examples this may be found by calculating the standard deviation of each row and column of the captured image 100, and finding the first and last discontinuities in the standard deviation of the rows and columns. In some examples finding the first and last discontinuities of the rows and columns may be performed using edge detection, for example a Canny edge detector.

5 This may correspond with the candidate boundary location at 220 in the method of figure 2a.

[0044] In some examples, the boundary may be an edge or a corner. In some examples a bounding box 110 may be defined based on at least one candidate location. The bounding box 110 may provide an estimate of the area occupied by the document 120. In some examples, the bounding box 110 may be defined based on at least four candidate locations 110a-110d, each candidate location being an estimate of the location of a document corner 120a-120d of the document 120.

[0045] In some examples the validity of the candidate location may be determined at 330 based on pixel values, around the candidate locations, of a first channel of the captured image 100, and pixel values of at least one other channel of the captured image 100, wherein the first channel is different from the at least one other channel. Carrying out the validation using a first channel and at least one other channel may lead to a more robust or accurate validation. The at least one other channel may have different information content compared with the first channel, such that the validation operation may make use of more of the information available in the captured image. In some examples, this additional use of information in the captured image does not significantly increase processing requirements.

[0046] In some examples, one or more of the first channel and at least one other channel are channels of the captured image. In some examples, one or more of the first channel and at least one other channel are not channels of the captured image and are derived from the captured image prior to receiving the image at operation 310, and are received prior to operation 330. In some examples, one or more of the first channel and at least one other channel are not channels of the captured image and are determined as part of method 300, prior to operation 330.

[0047] In some examples the first channel is obtained by grayscaling the captured image. In some examples the first channel is a luminance channel. In some examples the at least one other channel includes one or more of a red channel, a green channel and a blue channel. In some examples the at least one other channel includes all of a red channel, a green channel and a blue channel.

[0048] In some examples, the validity of the candidate location may be determined at 330 based on a determination of a first statistic of the pixel values around the candidate location of the first channel, and a comparison between the first statistic and a background statistic

associated with the background area, wherein the background area 130 is the area outside the document 130 (or, in some examples, the area outside the document assuming that the candidate location of the boundary is a boundary of the document).

[0049] In some examples, the first statistic is an indication of variation in the pixel values of the first channel across the candidate boundary. In some examples, the background statistic is an indication of variation in the pixel values (e.g. of the first channel) across the background. In some examples, where the variation across the candidate boundary is determined to be less than the variation of the background, it may be determined that the candidate boundary is not a valid boundary.

[0050] In some examples the first statistic is an average (1D, one-dimensional) standard deviation of pixel values in the first channel around the candidate location. For example, if the candidate location is a candidate edge that is approximately vertical (e.g. crossing more rows than columns per unit length of the edge) the standard deviation may be calculated for each row, based on a set of pixels around the candidate edge (e.g. 10 pixels either side of the candidate edge). The average (e.g. mean) value of the standard deviations associated with the candidate edge may then be calculated as the first statistic. Herein, this mean of the 1D standard deviation of pixel values around the candidate edge is referred to as stdmean.

[0051] In some examples an upper bound (TH) and a lower bound (TL) are defined based on the background statistic. The background statistic may be the standard deviation of pixel values corresponding to the background area 130. In some examples, upper and lower thresholds/bounds may be determined for a horizontal direction, to give TH_h and TL_h (high and low thresholds/bounds in the horizontal direction). Similarly, upper and lower thresholds/bounds may be determined for a vertical direction, to give TH_v and TL_v (high and low thresholds/bounds in the vertical direction). The horizontal bounds may be used when validating horizontal candidate boundaries, and the vertical bounds may be used when validating vertical candidate boundaries.

[0052] In some examples, the horizontal upper bound is the highest standard deviation among 1D standard deviations of pixel values in columns of background pixels. That is, the 1D standard deviation is determined for each column of a plurality of columns of background pixels, and the upper bound is the highest of these standard deviations. Similarly, the horizontal lower bound is the lowest standard deviation among 1D standard deviations of pixel values in columns of background pixels. Similarly, the vertical upper bound is the highest standard deviation among 1D standard deviations of pixel values in rows of background pixels and the vertical lower bound is the lowest standard deviation among 1D standard deviations of pixel values in rows of background pixels. In some examples, the background statistic is based

on pixel values of a single channel. In some examples, the single channel is the same as the first channel. In some examples, the single channel is a luminance channel. In some examples, the background statistic may be calculated based on an image of the background without the document (e.g. an image of the background captured before the captured image including the document).

[0053] In some examples the candidate location is determined to be not valid if the first statistic is within the range between a lower bound, TL, and the upper bound, TH. For example, the candidate location is determined to be not validated if $TL \leq \text{stdmean} \leq TH$.

[0054] In some examples the background statistic may be previously determined for a particular background (e.g. a particular scanner backplate). Accordingly, in some examples the background statistic is not calculated for each captured image. In some examples, the background statistic may be calculated based on the background area of the captured image 100, and may be calculated for each successive captured image.

[0055] In some examples, the validation may be based on at least one other statistic. In some examples, the validity of the candidate location may be based on a comparison of the at least one other statistic to at least one other (or further) background statistic associated with the background area. In some examples, the validation may be based on a difference between a value of the at least one other statistic and a value of the at least one other background statistic. The difference may be indicative of a difference between image properties (e.g. color) before and after the boundary.

[0056] In some examples, the at least one other statistic may be based on a median of pixel values around the candidate location. In some examples the at least one other channel is different from the first channel. For example, if the first channel is a grayscale channel, the at least one other channel may be one or more channels other than the grayscale channel.

[0057] In some examples the at least one other channel may be a red channel, a green channel and a blue channel. captured image 100 may be a RGB image, comprising a red channel, a green channel and a blue channel.

[0058] The at least one other statistic may include respective median values for each of the red, green and blue channels. For example, a median of the pixel values in the red color channel in a region of the candidate boundary on the document side of the candidate boundary may be determined, and referred to herein as median_r(doc). The region of the candidate boundary, over which the median is determined, may be, for example, a region within the bounding box and within a predetermined number of pixels of the candidate boundary. For example, if the candidate boundary is a vertical edge, for each row of pixels

containing the edge, the region may be 10 pixels of that row that are within the bounding box and closest to the vertical edge. Other numbers of pixels may be used to define the region. For example, the region may be 5 pixels from the edge, 20 pixels from the edge, etc.

[0059] Similarly, median document values in the region of the candidate boundary may be determined for the green and blue channels, to give $\text{median_g}(\text{doc})$ and $\text{median_b}(\text{doc})$, respectively. In some examples the median may be calculated for channels other than red, green and blue channels.

[0060] In some examples the at least one other background statistic may include respective median values for each of the red, green and blue channels (or other channels corresponding with the channels used to determine the at least one other statistic). For example, a median of the pixel values in the red color channel in a region of the candidate boundary on the background side of the candidate boundary may be determined, and referred to herein as $\text{median_r}(\text{bg})$. The region of the candidate boundary, over which the median is determined, may be, for example, a region along the candidate boundary and outside the bounding box, within a predetermined number of pixels of the candidate boundary. For example, if the candidate boundary is a vertical edge, for each row of pixels containing the edge, the region may be 10 pixels of that row that are outside the bounding box and closest to the vertical edge. Other numbers of pixels may be used to define the region. For example, the region may be 5 pixels from the edge, 20 pixels from the edge, etc.

[0061] Similarly, median background values in the region of the candidate boundary may be determined for the green and blue channels, to give $\text{median_g}(\text{bg})$ and $\text{median_b}(\text{bg})$, respectively. In some examples the median may be calculated for channels other than red, green and blue channels.

[0062] A difference between the at least one other statistic and the at least one other background statistic may be evaluated. For example, diff_median_r may be evaluated as $|\text{median_r}(\text{doc}) - \text{median_r}(\text{bg})|$. Similarly, $\text{diff_median_g} = |\text{median_g}(\text{doc}) - \text{median_g}(\text{bg})|$ and $\text{diff_median_b} = |\text{median_b}(\text{doc}) - \text{median_b}(\text{bg})|$.

[0063] In some examples the candidate location may be determined to be not valid if $\text{diff_median_r} \leq \text{TH_color}$, and if $\text{diff_median_g} \leq \text{TH_color}$, and if $\text{diff_median_b} \leq \text{TH_color}$, where TH_color is a color contrast threshold. TH_color may be determined prior to performing the method. TH_color may be empirically found by, e.g. by a user or a device manufacturer. TH_color may be different for each color (or channel).

[0064] In some examples the validity of the candidate location may be determined based on both the determination of the first statistic of the pixel values around the candidate location

of the first channel and the determination of the at least one other statistic of the pixel values around the candidate location of the at least one other channel. In some examples, the candidate location may be determined to be valid upon meeting the criteria of:

1. (stdmean < TL or stdmean > TH); and
- 5 2. (diff_median_r > TH_color) or (diff_median_g > TH_color) or (diff_median_b > TH_color)

[0065] In some examples, one or more of the relational operators (less than, greater than) above may be replaced with the corresponding operator that allows for equality (less than or equal to, greater than or equal to). Other arrangements are possible for determining validity of
 10 a candidate location based on the first statistic, background statistic, at least one other statistic and at least one other background statistic.

[0066] In some examples, a new candidate location of the boundary may be found if the candidate location is determined to be invalid. This may comprise finding the next discontinuities of each row and column. In some examples a next discontinuity may be the
 15 next detected discontinuity along the same axis as the invalidated discontinuity in a direction that reduces the size of the bounding box. For example, if the discontinuity 225 of Figure 2a (i.e. the first discontinuity along the vertical direction) is found to be invalid, the next discontinuity is the next discontinuity that is arrived at moving down the page from the first discontinuity 225. The process may be repeated until a candidate boundary is found that
 20 satisfies the above requirements.

[0067] In some examples the candidate location may be one of a plurality of candidate locations, and validating the candidate location may further comprise finding the plurality of candidate locations, and determining the validity of each of the plurality of the candidate locations, wherein the plurality of candidate locations 110a-110d define a bounding box 110
 25 around the document.

[0068] In some examples, the first statistic and the background statistic may be based on one or more of smoothness, uniformity, entropy, gray level co-occurrence matrix, invariant moments, or spectral texture measures. In some examples, these may produce reliable results when the background has a 2D (two-dimensional) texture.

30 **[0069]** An example method 400 for validating a candidate location of a boundary between a document 120 and the background area 130 is described in accordance with figure 4. The method 400 may be used in operation 230 of figure 2a. In some examples, the method 400 may be used in other applications or used independently.

[0070] The captured image 100 may be received at 410. The captured image 100 may

include a document 120 and a background 130. A candidate location of a boundary between the document 120 and the background 130 may be determined 420 based on pixel values of a first channel of the captured image 100. The validity of the candidate location may be determined at 430, based on based on pixel values around the candidate location of at least a
5 second channel of the image, wherein the first channel is different from the second channel.

[0071] In some examples, at least one candidate location may be found at 420. In some examples this may be found by calculating the standard deviation of each row and column of a first channel of the captured image 100, and finding the first and last discontinuities in the standard deviation of the rows and columns. In some examples finding the first and last
10 discontinuities of the rows and columns may be performed using edge detection, for example a Canny edge detector. This may correspond with the candidate boundary location at 220 in the method of figure 2a.

[0072] In some examples the first channel is a luminance channel. In some examples, the first channel is derived from the captured image by grayscaling the captured image.

[0073] In some examples, the boundary may be an edge or a corner. In some examples a bounding box 110 may be defined based on at least one candidate location. The bounding box 110 may provide an estimate of the area occupied by the document 120. In some examples, the bounding box 110 may be defined based on at least four candidate locations 110a-110d, each candidate location being an estimate of the location of a document corner 120a-120d of
20 the document 120.

[0074] In some examples the validity of the candidate location may be determined at 430 based on pixel values, around the candidate locations, of a second channel of the captured image 100, wherein the first channel is different from the second channel. Carrying out the validation using at least a different channel from the channel used to detect the candidate
25 location may lead to improved accuracy. For example, using different channels may make use of more of the information contained in the captured image, and may lead to improved independence between the detection and validation of the candidate location. In some examples, the additional use of information in the captured image does not significantly increase processing requirements.

[0075] In some examples, the validity of the candidate location may be determined 430 based on a determination a first statistical measure of the pixel values around the candidate location of the first channel, and a comparison of the first statistical measure to a statistical measure associated with the background area.

[0076] In some examples, a plurality of second channels, each different from the first

channel, may be used to determine the validity of the candidate location.

[0077] In some examples, the first channel and/or the second channel are channels of the captured image. In some examples, the first channel and/or the second channel are not channels of the captured image and are derived from the captured image prior to receiving the image at operation 410, and are received prior to operation 430. In some examples, the first channel and/or the second channel are not channels of the captured image and are determined as part of method 400, prior to operation 430.

[0078] In some examples the first channel is obtained by grayscaling the captured image. In some examples the first channel is a luminance channel. In some examples the second channel is a red channel, a green channel or a blue channel. In some examples, there are three second channels, a red channel, a green channel or a blue channel, respectively.

[0079] In some examples, the validity of the candidate location may be determined based on a statistical measure. The statistical measure may be similar to the at least one other statistic of the method 300.

[0080] In some examples the candidate location may be determined to be not valid if $\text{diff_median_r} \leq \text{TH_color}$, and if $\text{diff_median_g} \leq \text{TH_color}$, and if $\text{diff_median_b} \leq \text{TH_color}$, where TH_color is a color contrast threshold, where diff_median_r, diff_median_g, diff_median_b and TH_color are as described in relation to the method 300.

[0081] Further, in some examples the candidate location may be determined to be valid upon meeting the criteria of:

1. ($\text{stdmean} < \text{TL}$ or $\text{stdmean} > \text{TH}$); and
2. ($\text{diff_median_r} > \text{TH_color}$) or ($\text{diff_median_g} > \text{TH_color}$) or ($\text{diff_median_b} > \text{TH_color}$),

Where stdmean, TL and TH are as described above in relation to method 300.

[0082] In some examples, stdmean, TL and TH are determined based on a single channel, and the single channel is the first channel (i.e. the channel used in operation 420).

[0083] In some examples, one or more of the relational operators (less than, greater than) above may be replaced with the corresponding operator that allows for equality (less than or equal to, greater than or equal to).

[0084] In some examples, a new candidate location of the boundary may be found if the candidate location is determined to be invalid. This may comprise finding the next discontinuities of each row and column. The validation may then be performed on the new candidate location.

[0085] In some examples the candidate location may be one of a plurality of candidate locations, and validating the candidate location may further comprise finding the plurality of candidate locations, and determining the validity of each of the plurality of the candidate locations, wherein the plurality of candidate locations 110a-110d define a bounding box 110 around the document.

[0086] Figure 5 illustrates an example method 500 to refine an estimate of a position of a document corner 120a-120d of a document 120 according to some examples. The method 500 may be practiced as a part of method 200 (e.g. in operation 260), or independently from method 200.

[0087] Method 500 may be applied to an image, such as a captured image 100 of a document 120 and a background 130. The document is inside a captured image 100, and is smaller than the captured image 100. The background 130 may correspond to the portion of the captured image 100 that is not the document 120.

[0088] At 510 a first estimated position of a corner 120a-120d of a document 120 is identified. This may be carried out as described above in relation to operations 220 to 240 of method 200, for example.

[0089] At 520 a window is defined around the first estimated position, wherein the area of the window is smaller than the area of the image. In some examples, this may be carried out as described in relation to operation 250 of method 200.

[0090] At 530, a second estimated position of the corner 120a-120d of the document 120 based on a portion of the image inside the window is identified. In some examples, this may correspond with operation 260 of method 200.

[0091] An example of the method 500 is described in more detail with reference to figure 6. In some examples a bounding box is defined around document 120, the bounding box may identify a first estimated position of the document corner 120a. in some examples the first estimated position may coincide with the bounding box corner 110a.

[0092] The first estimate may be identified based on at least one discontinuity in a statistical measure of pixel values of a plurality of rows and columns of the image. In some examples the statistical measure may be a 1D standard deviation of a channel of (or derived from) the captured image 100. In some examples, the channel is a luminance channel.

[0093] At 520 a window 610 is defined at or around the first estimated position. The corner of the window 610 may be aligned with the corner of the bounding box 110 such that the window is entirely inside the bounding box 110. In some examples, the window 610 may extend outside of the bounding box 110, such that the corner of the bounding box 110 is within the

window 610. The window 610 may have an area smaller than the bounding box 110. The window 610 may have a square or rectangular shape.

[0094] At 530 a second estimated position 620 of the corner 120a of the document 120 may be determined based on a portion of the captured image 100 inside the window 610. In some examples the second estimated position 620 is identified based on at least one discontinuity in a statistical measure of pixel values. In some examples, the pixel values include pixel values (e.g. an intensity of pixels) in a channel of the captured image. In some examples, the channel is a luminance channel. The statistical measure may be a 1D standard deviation along rows and/or columns. For example, a 1D standard deviation may be calculated for each row inside the window, and a discontinuity in the standard deviation may be detected. Similarly, a 1D standard deviation may be calculated for each column inside the window, and a discontinuity in the standard deviation may be detected.

[0095] In the refinement process of some examples, each window may include a single corner 120 of the document 110, such that a single discontinuity in the standard deviation is to be found for the rows, and a single discontinuity in the standard deviation is to be found for columns within each window. Where multiple discontinuities are detected, a discontinuity may be selected as the second estimated position 620. The second estimated position 620 may be based on the discontinuity closest to the corner of the window corresponding to the bounding box corner that is being refined (or the corner of the document that is being approximated). For example, if the window is at (i.e. approximating) the top right corner of the document, such that the top right bounding box corner is being refined, the rightmost discontinuity may be selected in the horizontal direction and topmost discontinuity may be selected in the vertical direction.

[0096] As shown in figure 6, by performing a second refinement of the estimated position of the corner 120a of the document 120 inside the window an improved accuracy may be obtained for the estimate of the corner location.

[0097] In some examples, further processing may be carried out within the refined window to identify the document of the corner more accurately. In some examples a corner detection method may be applied inside refined the window. For example, a corner detection method may be applied as described in relation to operation 270 of method 200. In some examples, refining the window/estimated position may lead to a reduced window size/area. This may reduce a processing load associated with performing a corner detection method in the window.

[0098] In some examples the location of the corner of the document 120 may be estimated based on a comparison of the second estimated position 620 and possible corner locations identified by the corner detection method. In some examples the possible corner locations in the window identified by the corner detection method may be a subset of corner locations

initially identified in the window by the corner detection method. For example, a strength value associated with corners detected by the corner detection method may be used to obtain the subset of corner locations.

[0099] In some examples the corner detection method is based on a determination of at least a partial derivative of the intensity of pixels inside the window. In some examples the corner detection method is based on at least one of a Harris corner detector, a SUSAN corner detector, and a minimum eigenvalue corner detector.

[00100] In some examples determining the first subset of possible corner locations comprises: determining a strength of each possible corner location and selecting the first subset based on the strength.

[00101] In some examples approximating the location of the corner of the page comprises comparing the distance from each of the first subset of possible corner locations to the second estimate 620 the location of the corner 120a.

[00102] An image processing apparatus 700 according to some examples is described in relation to figure 7a. The apparatus 700 comprises image receiving circuitry 710, document location circuitry 720, and validation circuitry 730.

[00103] The image receiving circuitry 710 is configured to receive an image. The image may be a captured image, and may include a document and a background.

[00104] The document location circuitry 720 is configured to find a candidate location of a boundary between the document and the background based on pixel values of a first channel of the image. The document location circuitry 720 may be arranged to operate as described above in relation to operation 220 of Figure 2a, operation 320 of Figure 3 or operation 420 of Figure 4.

[00105] The validation circuitry 730 is configured to validate the candidate location of the boundary, as described above in relation to operation 230 of Figure 2, operation 330 of Figure 3 or operation 430 of Figure 4.

[00106] An image processing apparatus 750 according to some examples is described in relation to figure 7b. The apparatus 750 comprises first corner estimation circuitry 760, windowing circuitry 770, and second corner estimation circuitry 780.

[00107] The first corner estimation circuitry 760 is configured to receive an image. The image may be a captured image, and may include a document and a background. The first corner estimation circuitry 760 is arranged to determine a first estimated position a corner of the document in the image. The first corner estimation circuitry 760 may be arranged to operate as

described above in relation to operation 240 of Figure 2a or operation 510 of Figure 5.

[00108] The windowing circuitry 770 is configured to define a window around the first estimated position, with the area of the window being smaller than the area of the image. The windowing circuitry 770 may be arranged to operate as described above in relation to operation 250 of Figure 2a or operation 520 of Figure 5.

[00109] The second corner estimation circuitry 780 is arranged to identify a second estimated position of the corner of the document based on a portion of the image inside the window. The second corner estimation circuitry 780 may be arranged to operate as described above in relation to operation 260 of Figure 2a or operation 530 of Figure 5.

[00110] Figure 8 is a schematic drawing of an example of a system 800 to automatically detect corners of an image of a captured image 100.

[00111] System 800 may include a processor 840. Processor 840 may include one or more intercommunicating processors that may operate in accordance with programmed instructions. For example, at least a component of processor 840 may be incorporated into a processor of a computer or into a processor or controller of an image scanner. All such processors and controllers are considered as incorporated into processor 840.

[00112] Processor 840 may communicate with memory 850. Memory 850 may include one or more volatile or non-volatile memory devices capable of storing data. For example, memory 850 may be utilized to store, among other data, programmed instructions to operate processor 840, acquired scanned images, results generated during or as a result of processing scanned images, or parameters or data utilized in processing scanned images.

[00113] Processor 840 may communicate with a data storage device 860. Data storage device 860 may represent a computer readable medium in the form of one or more fixed or removable non-volatile data storage devices. For example, data storage device 860 may be utilized to store, among other data, programmed instructions to operate processor 840, acquired scanned images, results generated during or as a result of processing scanned images, or processing scanned images.

[00114] Processor 840 may communicate with an imaging device represented by optical scanner sensor 870. Optical scanner sensor 870 may include a sensor, or a collection or array of sensors, such as a camera or scanner head, that are capable of acquiring a digital image of a document 120, or other object. For example, processor 840 may control scanner sensor 870 to acquire a captured image 100, the captured image including a document 120 and a background 130. The image may be in the form of an array (e.g. rows and columns) of pixels, each pixel including image data (e.g. polychromatic RGB data). Processor 840 may process

the acquired image or store the acquired image in memory 850 or data storage device 860.

[00115] For example, processor 840 may control scanner sensor 870 to acquire an image of document 120. Document 120 may be scanned when placed on or held to the surface of the optical scanner sensor 870. For example, the surface may be a conveyer belt or other transport device for transporting document 120 to or past scanner sensor 870. For example, operation of a conveyance device represented the surface may be controllable by processor 840.

[00116] In some examples, the processor 840 may be arranged to operate as the image receiving circuitry 710, document location circuitry 720 and validation circuitry 730 of Figure 7a. For example, memory 850 and/or data storage device 860 may include program instructions that, when executed, cause the processor to operate as the image receiving circuitry 710, document location circuitry 720 and validation circuitry 730 of Figure 7a.

[00117] In some examples, the processor 840 may be arranged to operate as the first corner estimation circuitry 760, windowing circuitry 770, and second corner estimation circuitry 780 of Figure 7b. For example, memory 850 and/or data storage device 860 may include program instructions that, when executed, cause the processor to operate as the first corner estimation circuitry 760, windowing circuitry 770, and second corner estimation circuitry 780 of Figure 7b.

[00118] Figure 9a illustrates a computer readable medium 900 according to some examples. The computer readable medium 900 stores units, with each unit including instructions that, when executed, cause a processor 840 or other processing device to perform particular operations. The computer readable medium 900 includes a receiving unit 910 including instructions that, when executed, cause a processing device 840 to receive an image comprising a document and a background area. The computer readable medium 900 also includes a boundary finding unit 920 including instructions that, when executed, cause the processing device 840 to find a candidate location of a boundary of the document. The computer readable medium 900 also includes a boundary validation unit 930 including instructions that, when executed, cause the processing device 840 to validate the candidate location of the boundary of the document. In some examples the receiving unit 910 may cause the processor to operate in accordance with operation 310 of method 300 or operation 410 of method 400. In some examples the boundary finding unit 920 and boundary validation unit 930 may cause the processor to operate in accordance with operations 220 and 230 of method 200, operations 320 and 330 of method 300, or operations 420 and 430 of method 400. The units of the computer readable medium 900 may cause a processing device 840 to operate in accordance with any of the examples described herein.

[00119] Figure 9b illustrates a computer readable medium 950 according to some examples.

The computer readable medium 950 stores units, with each unit including instructions that, when executed, cause a processor 840 or other processing device to perform particular operations. The computer readable medium 950 includes a first corner estimating unit 960 including instructions that, when executed, cause a processing device 840 to determine a first
5 estimated position of a corner of a document in a captured image. The computer readable medium 950 also includes a windowing unit 970 including instructions that, when executed, cause the processing device 840 to define a window around the first estimated position. The computer readable medium 950 also includes a second corner estimation unit 980 including
10 instructions that, when executed, cause the processing device 840 to identify a second estimated position of the corner of the document. In some examples the first corner estimating unit 960 may cause the processor to operate in accordance with operation 240 of method 200 or operation 510 of method 500. In some examples the windowing unit 970 may cause the
processor to operate in accordance with operation 250 of method 200 or operation 520 of method 500. In some examples the second corner estimation unit 980 may cause the
15 processor to operate in accordance with operation 260 of method 200 or operation 530 of method 500.

[00120] The units of the computer readable medium 900 may cause a processing device 840 to operate in accordance with any of the examples described herein.

[00121] Figure 10 illustrates a printing device 1000 according to some examples. The
20 arrangement of Figure 10 is a printing device for applying a 2-dimensional (2D) image to a substrate. The printed output 1040 may be a printed medium that is printed by the printing device 1000.

[00122] Print instructions 1005, such as a print job, are received at input 1010. Print instructions 1005 may include a digital description of a document to be printed on a substrate.

[00123] The printing device 1000 may include one or more control elements, illustrated as
25 controller 1020, to control the various components of the printing device 1000. The controller 1020 may be implemented in software, hardware, firmware or a combination of these.

[00124] The controller 1020 controls an image fixing section 1030 to print or otherwise fix the document to a substrate to generate printed output 1040. The image fixing section 1030 may
30 be, for example, a digital press, offset press, inkjet printer, toner printer, etc.

[00125] The printed output 1040 is scanned by scanner 1050, and may subsequently be output from the printing device 1000.

[00126] The scanner 1050 outputs a digital representation (a captured image) 1055 of the printed output 1040 to the document detection circuitry. The document detection circuitry may

operate in accordance with any example herein to determine a boundary of a document 120 in the captured image 100.

[00127] Comparison circuitry 1070 may receive a reference document 1025 from controller 220, based on the document to be printed according to the print instructions 1005. The

5 comparison circuitry 1070 may compare the reference document 1025 with the document 120 in the captured image 100, using the determined boundary of the document 120 in the captured image 100 to align or register the document 120 in the captured image 100 with the reference document 1025.

[00128] The comparison circuitry 1070 may determine whether or not there are any defects in
10 the printed output 1040, based on the comparison between the reference image 1025 and the document 120 in the captured image 100.

[00129] In some examples a color space conversion may be performed on one or more of the document in the print instructions 1005, the captured image 1055 and/or the reference image. For example, the document in the print instructions 1005 may be received as, or converted to,
15 a CMYK color space. The captured image may be generated by the scanner in a RGB color space. The reference image may be in the CMYK color space, RGB color space or some other color space. The reference image and the captured image may be converted by the comparison circuitry, or prior to arriving at the comparison circuitry 1070, to the same color space to facilitate the comparison.

20 **[00130]** In some examples, the scanner 1050 may be an in-line scanner.

[00131] In figure 10, the various components are illustrated as being part of the printing device 1000. However, in some examples, some components may be provided separately from the printing device, in different functional processing units, for example.

[00132] In some examples the printing device 1000 may be a press, such as a commercial
25 press. In some examples, in order to maintain high throughput, the printed output may be moving at high speed (e.g. around 1280 mm/s). In some examples, the document detection circuitry 1060 and comparison circuitry 1070 may be arranged to operate at a speed that is consistent with the speed of the printing process. Efficient processing by the document detection circuitry 1060 and the comparison circuitry 1070 may assist in operating at a speed
30 consistent with the printing process while avoiding a significant cost increase associated with hardware having greater processing capability.

[00133] In some examples, the handling of the printed output 1040 may be dependent on whether or not any defects are determined to be present by the comparison circuitry 1070. For example, if a defect is determined to be present, the output may be paused until a user has

intervened, or an output path of the printed output 1040 may be changed.

[00134] In some examples of printing device 1000, the printed output 1040 may be moved and controlled by medium handling elements, such as rollers, conveyor belts, etc. In some examples, a maximum expected skew of the document 120 in the captured image 100 may be predicted based on the medium handling elements. For example, in some cases a skew in the printed output 1040 exceeding a threshold value (e.g. of around 3 degrees) may lead to the medium jamming. In such a case, the skew of the document 120 in the captured image 100 may be expected to be less than the threshold value (e.g. 3 degrees). In some examples, parameters of the document detection/boundary detection method may be set based on the threshold (maximum expected) skew. For example, the sizes of the windows in one or more of operations 250, 260, 520 and 530 of methods 200 and 500 may be set based on the threshold (maximum expected) skew.

[00135] may be set based on the assumption that the skew of the document in the captured image is less than the threshold skew.

[00136] Examples herein refer to a captured image. The method of producing the image is not particularly limited, and the image may be received (as a received image) from any image source that produces suitable images.

[00137] Throughout the description and claims of this specification, the words “comprise” and “contain” and variations of them mean “including but not limited to”, and they are not intended to (and do not) exclude other components, integers or operations. Throughout the description and claims of this specification, the singular encompasses the plural unless the context implies otherwise. In particular, where the indefinite article is used, the specification is to be understood as contemplating plurality as well as singularity, unless the context implies otherwise.

[00138] Features, integers or characteristics described in conjunction with a particular aspect or example are to be understood to be applicable to any other aspect or example described herein unless incompatible therewith. All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the operations of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or operations are mutually exclusive. Examples are not restricted to the details of any foregoing examples. The Examples may extend to any novel one, or any novel combination, of the features disclosed in this specification (including any accompanying claims, abstract and drawings), or to any novel one, or any novel combination, of the operations of any method or process so disclosed.

[00139] Examples provide a method comprising: receiving an image comprising a document and background area; finding a candidate location of a boundary between the document and the background area; and checking a validity of the candidate location of the boundary, based on pixel values around the candidate location of a first channel of the image, and pixel values around the candidate location of at least one other channel of the image, wherein the at least one other channel is different from the first channel.

[00140] In some examples, checking the validity of the candidate location of the boundary is based on at least one of: determining a first statistic of the pixel values around the candidate location of the first channel, and comparing the first statistic to a background statistic associated with the background area; and determining at least one other statistic of the pixel values around the candidate location of the at least one other channel, and comparing the at least one other statistic to at least one other background statistic associated with the background area.

[00141] In some examples, the first statistic is based on a luminance of the pixel values around the candidate location.

[00142] Some examples further comprise finding a new candidate location of the boundary if the candidate location is determined to be invalid.

[00143] In some examples, the boundary is a corner or an edge.

[00144] In some examples, wherein the candidate location is one of a plurality of candidate locations, and the method further comprises finding the plurality of candidate locations, and checking the validity of each of the plurality of the candidate locations, wherein the plurality of candidate locations defines a bounding box around the document.

[00145] In some examples, the candidate location of the boundary between the document and the background area is found based on a third channel.

[00146] In some examples, the third channel is the same as the first channel.

[00147] Examples provide a method comprising: receiving an image comprising a document and background area; finding a candidate location of a boundary between the document and the background area based on pixel values of a first channel; and checking a validity of the candidate location of the boundary, based on pixel values

around the candidate location of at least a second channel of the image, wherein the first channel is different from the second channel.

[00148] In some examples, checking the validity of the candidate location of the boundary is based on at least one of: determining a first statistical measure of the pixel values around the candidate location of the first channel, and comparing the first statistical measure to a background statistical measure associated with the background area, and determining at least one other statistical measure of the pixel values around the candidate location of the at least second channel, and comparing the at least one other statistical measure to at least one other background statistic associated with the background area.

[00149] In some examples, the validity of the candidate location of the boundary is checked based on a third channel.

[00150] In some examples, a third channel is the same as the first channel.

[00151] Examples provide an apparatus comprising: image receiving circuitry to receive an image, the image comprising a document and background area; document location circuitry to find a candidate location of a boundary between the document and the background area based on pixel values of a first channel; and validation circuitry to validate the candidate location of the boundary, based on pixel values around the candidate location of at least a second channel of the image, wherein the first channel is different from the second channel.

[00152] In some examples, the validation circuitry is to perform at least one of: determine a first statistical measure of the pixel values around the candidate location of the first channel, and compare the first statistical measure to a background statistical measure associated with the background area; and determine at least one other statistical measure of the pixel values around the candidate location of the at least second channel, and comparing the at least one other statistical measure to at least one other background statistic associated with the background area.

[00153] Examples provide a non-transitory computer readable medium having stored thereon instructions that when executed, cause a processor to perform the method according to some examples.

[00154] Examples provide a method comprising: identifying a first estimated position of

a corner of a document, the document inside an image; defining a window around the first estimated position, wherein the area of the window is smaller than the area of the image; identifying a second estimated position of the corner of the document based on a portion of the image inside the window.

- 5 **[00155]** In some examples, the first estimate is identified based on at least one discontinuity in a statistical measure of intensity of a plurality of rows and columns of the image.

- 10 **[00156]** In some examples, the second estimated position is identified based on at least one discontinuity in a statistical measure of intensity of a plurality of rows and columns of the portion of the image inside the window.

[00157] Some examples further comprise: applying a corner detection method inside the window, resulting in a plurality of possible corner locations; and determining a first subset of potential corner locations from the plurality of potential corner locations.

- 15 **[00158]** Some examples further comprise: defining a location of the corner of the document based on a comparison of the second estimated position and the first subset of possible corner locations.

[00159] In some examples, the corner detection method is based on a determination of at least a partial derivative of the intensity of pixels inside the window.

- 20 **[00160]** In some examples, the corner detection method is based on at least one of a Harris corner detector, a SUSAN corner detector, and a minimum eigenvalue corner detector.

[00161] In some examples, determining the first subset of possible corner locations comprises: determining a strength of each possible corner location and selecting the first subset based on the strength.

- 25 **[00162]** In some examples, defining the location of the corner of the page comprises comparing the distance from each of the first subset of possible corner locations to the second estimate the location of the corner.

- 30 **[00163]** Some examples comprise an apparatus comprising: first corner estimate circuitry to determine a first estimated position of a corner of a document, the document inside a image; windowing circuitry to define a window around the first estimated

position, wherein the area of the window is smaller than the area of the image; and second corner estimate circuitry to identify a second estimated position of the corner of the document based on a portion of the image inside the window.

[00164] Some examples further comprise corner detection circuitry to apply a corner detection method inside the window, resulting in a plurality of possible corner locations; and to determine a first subset of potential corner locations from the plurality of potential corner locations.

[00165] In some examples, the corner detection method is based on a determination of at least a partial derivative of the intensity of pixels inside the window.

[00166] In some examples, the corner detection circuitry is to determine the first subset of possible corner locations by determining a strength of each possible corner location and selecting the first subset based on the strength.

[00167] In some examples, the apparatus further comprises: document corner defining circuitry to define the location of the corner of the page based on a comparison the distance from each of the first subset of possible corner locations to the second estimate the location of the corner.

[00168] Some examples comprise a non-transitory computer readable medium having stored thereon instructions that when executed, cause a processor to perform the method of some examples.

CLAIMS

5

1. A method comprising:

receiving an image comprising a document and background area;

finding a candidate location of a boundary between the document and the background area; and

10

checking a validity of the candidate location of the boundary, based on pixel values around the candidate location of a first channel of the image, and pixel values around the candidate location of at least one other channel of the image, wherein the at least one other channel is different from the first channel.

15

2. The method according to claim 1, wherein checking the validity of the candidate location of the boundary is based on at least one of:

determining a first statistic of the pixel values around the candidate location of the first channel, and comparing the first statistic to a background statistic associated with the background area; and

20

determining at least one other statistic of the pixel values around the candidate location of the at least one other channel, and comparing the at least one other statistic to at least one other background statistic associated with the background area.

25

3. The method according to claim 2, wherein the first statistic is based on a luminance of the pixel values around the candidate location.

4. The method according to claim 1, further comprising finding a new candidate location of the boundary if the candidate location is determined to be invalid.

30

5. The method according to claim 1, wherein the boundary is a corner or an edge.

6. The method according to claim 1, wherein the candidate location is one of a plurality of candidate locations, and the method further comprises finding the plurality of candidate locations, and checking the validity of each of the plurality of the candidate locations,

wherein the plurality of candidate locations defines a bounding box around the document.

7. The method according to claim 1, wherein the candidate location of the boundary
5 between the document and the background area is found based on a third channel.
8. The method according to claim 7, wherein the third channel is the same as the first channel.
- 10 9. A method comprising:
receiving an image comprising a document and background area;
finding a candidate location of a boundary between the document and the
background area based on pixel values of a first channel; and
checking a validity of the candidate location of the boundary, based on pixel
15 values around the candidate location of at least a second channel of the image,
wherein the first channel is different from the second channel.
10. The method according to claim 9, wherein checking the validity of the candidate
location of the boundary is based on at least one of:
20 determining a first statistical measure of the pixel values around the candidate
location of the first channel, and comparing the first statistical measure to a background
statistical measure associated with the background area, and
determining at least one other statistical measure of the pixel values around the
candidate location of the at least second channel, and comparing the at least one other
25 statistical measure to at least one other background statistic associated with the
background area.
11. The method according to claim 9, wherein the validity of the candidate location of the
boundary is checked based on a third channel.
30
12. The method according to claim 11, wherein the third channel is the same as the first
channel.
13. An apparatus comprising:
35 image receiving circuitry to receive an image, the image comprising a document
and background area;

document location circuitry to find a candidate location of a boundary between the document and the background area based on pixel values of a first channel; and

validation circuitry to validate the candidate location of the boundary, based on pixel values around the candidate location of at least a second channel of the image, wherein the first channel is different from the second channel.

14. The apparatus of claim 13, wherein the validation circuitry is to perform at least one of:
- determine a first statistical measure of the pixel values around the candidate location of the first channel, and compare the first statistical measure to a background statistical measure associated with the background area; and
 - determine at least one other statistical measure of the pixel values around the candidate location of the at least second channel, and comparing the at least one other statistical measure to at least one other background statistic associated with the background area.
15. A non-transitory computer readable medium having stored thereon instructions that when executed, cause a processor to perform the method of claim 1.

1/12




Figure 1






Figure 2a

3/12




230-295 of method 200

Figure 2b


4/12

300



310

320



330




Figure 3

5/12

400





Figure 4


6/12

500



510

520



530




Figure 5




Figure 6

8/12




Figure 7a

9/12




Figure 7b




Figure 8

11/12




Figure 9a




Figure 9b

12/12




Figure 10

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2017/075010

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06K9/00
 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)

G06K G06T H04N

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, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2001/014183 A1 (SANSOM-WAI CINDY Y [US] ET AL) 16 August 2001 (2001-08-16)	1,7-9, 11-13,15
Y	abstract; figure all paragraphs [0013], [0072], [0073], [0076], [0077], [0079] -----	2-6,10, 14
Y	US 2015/324639 A1 (MACCIOLA ANTHONY [US]) 12 November 2015 (2015-11-12) abstract; figure all paragraphs [0014], [0015], [0092], [0093], [0097], [0099], [0107], [0111], [0113], [0142] - [0145] ----- -/-	2,10,14



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

14 November 2017

Date of mailing of the international search report

22/11/2017

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

Winne, Dominique

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2017/075010

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 6 956 587 B1 (ANSON DAVID L A [US]) 18 October 2005 (2005-10-18) abstract; figure all column 1, line 55 - line 60 column 5, line 40 - line 45 column 6, line 10 - line 50 column 7, line 25 - line 40 -----	3-6

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2017/075010

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2001014183	A1	16-08-2001	EP 0949802 A2 13-10-1999
			JP H11313207 A 09-11-1999
			US 6430320 B1 06-08-2002
			US 2001014183 A1 16-08-2001

US 2015324639	A1	12-11-2015	NONE

US 6956587	B1	18-10-2005	NONE
