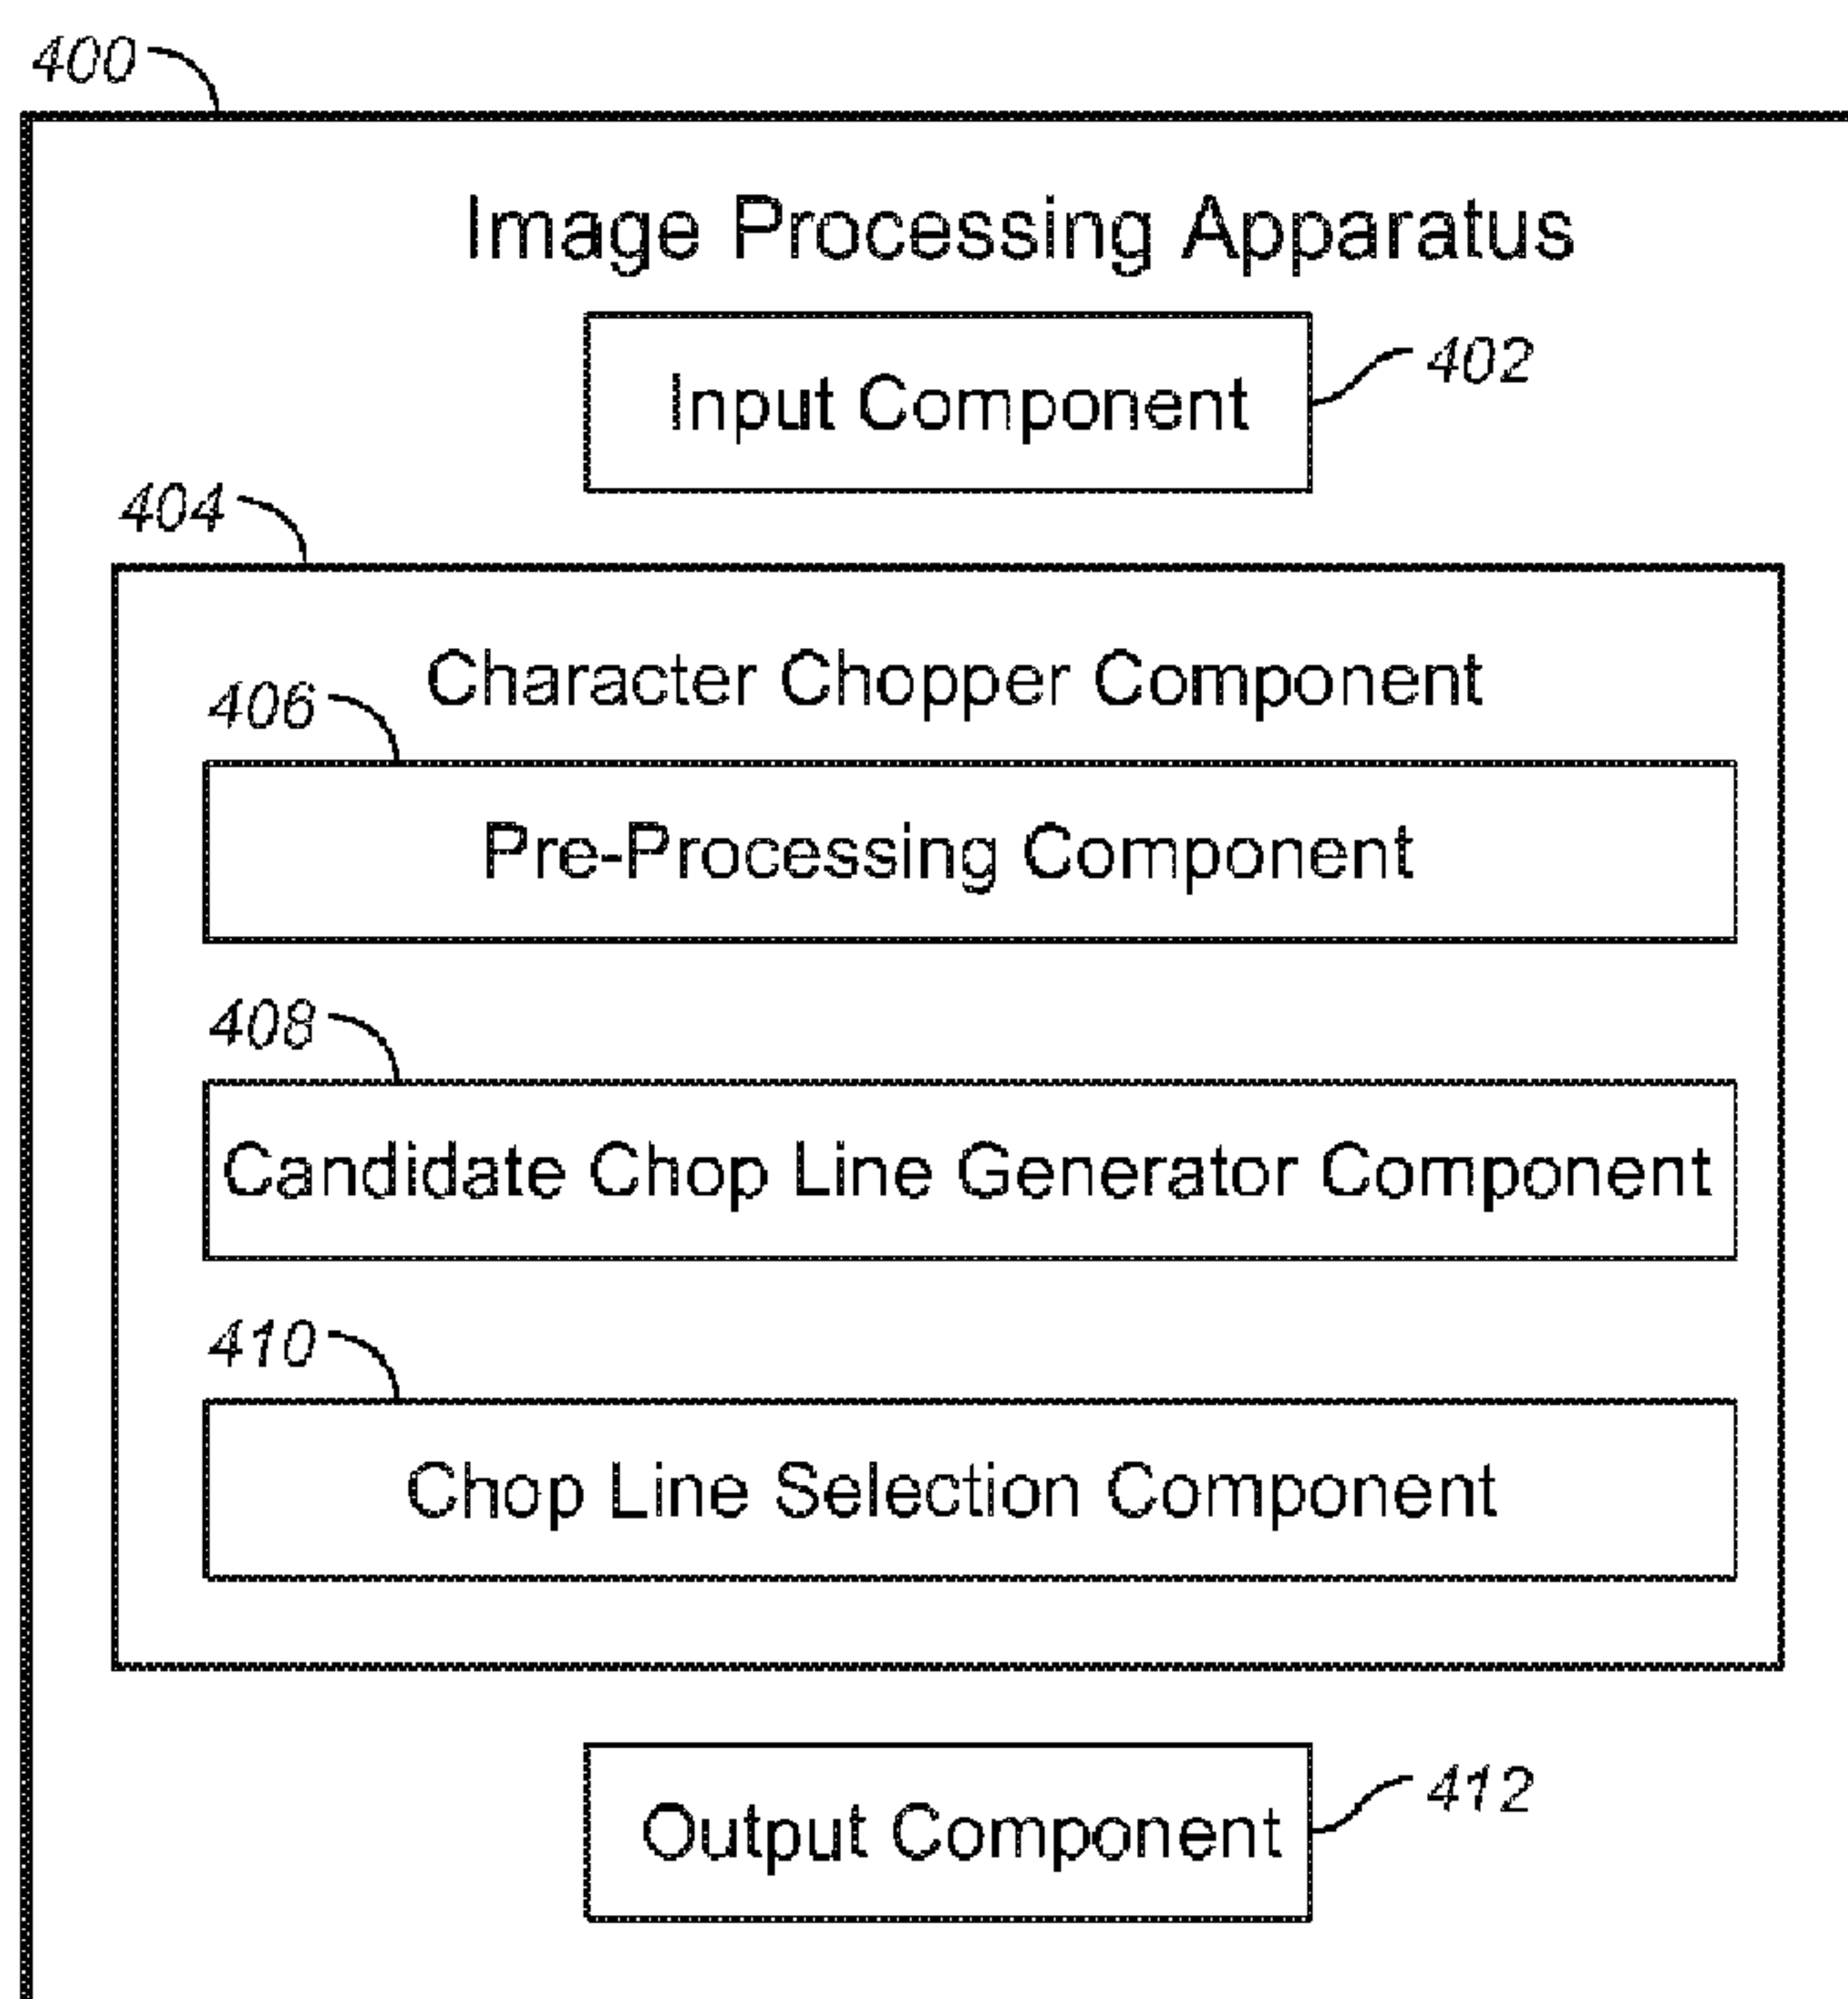




(86) Date de dépôt PCT/PCT Filing Date: 2011/04/28
 (87) Date publication PCT/PCT Publication Date: 2011/11/17
 (45) Date de délivrance/Issue Date: 2017/07/04
 (85) Entrée phase nationale/National Entry: 2012/10/24
 (86) N° demande PCT/PCT Application No.: US 2011/034242
 (87) N° publication PCT/PCT Publication No.: 2011/142977
 (30) Priorité/Priority: 2010/05/10 (US12/776,576)

(51) Cl.Int./Int.Cl. *H04N 1/387* (2006.01),
G06K 9/00 (2006.01), *H04N 1/00* (2006.01),
H04N 1/04 (2006.01)
 (72) Inventeur/Inventor:
NIJEMCEVIC, DJORDJE, US
 (73) Propriétaire/Owner:
MICROSOFT TECHNOLOGY LICENSING, LLC, US
 (74) Agent: SMART & BIGGAR

(54) Titre : SEGMENTATION D'UNE REPRESENTATION BINAIRE DE MOTS EN DES CARACTERES OU GLYPHES INDIVIDUELS PENDANT UN PROCESSUS OCR
 (54) Title: SEGMENTATION OF A WORD BITMAP INTO INDIVIDUAL CHARACTERS OR GLYPHS DURING AN OCR PROCESS



(57) **Abrégé/Abstract:**

An image processing apparatus is provided that includes a character chopper component that segments words into individual characters in a bitmap of a textual image undergoing an OCR process. The Character chopper component is configured to produce a set of (possibly curved) chop-lines which divide a bitmap of any given word into its individual character or glyph candidates. Cases where an input bitmap contains two separate words are handled by marking a place where those words should be split. The character segmentation algorithm computes the set of vertically oriented, curved chop-lines by considering glyph and background colors in a given word bitmap. The set is filtered afterwards using various heuristics, in order to preserve those lines that indeed do separate a word's glyphs and minimize the number of those that do not.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
17 November 2011 (17.11.2011)(10) International Publication Number
WO 2011/142977 A3

(51) International Patent Classification:

H04N 1/387 (2006.01) H04N 1/00 (2006.01)
H04N 1/04 (2006.01) G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2011/034242

(22) International Filing Date:

28 April 2011 (28.04.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

12/776,576 10 May 2010 (10.05.2010) US

(71) Applicant (for all designated States except US): MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventor: NIJEMCEVIC, Djordje; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

(54) Title: SEGMENTATION OF A WORD BITMAP INTO INDIVIDUAL CHARACTERS OR GLYPHS DURING AN OCR PROCESS

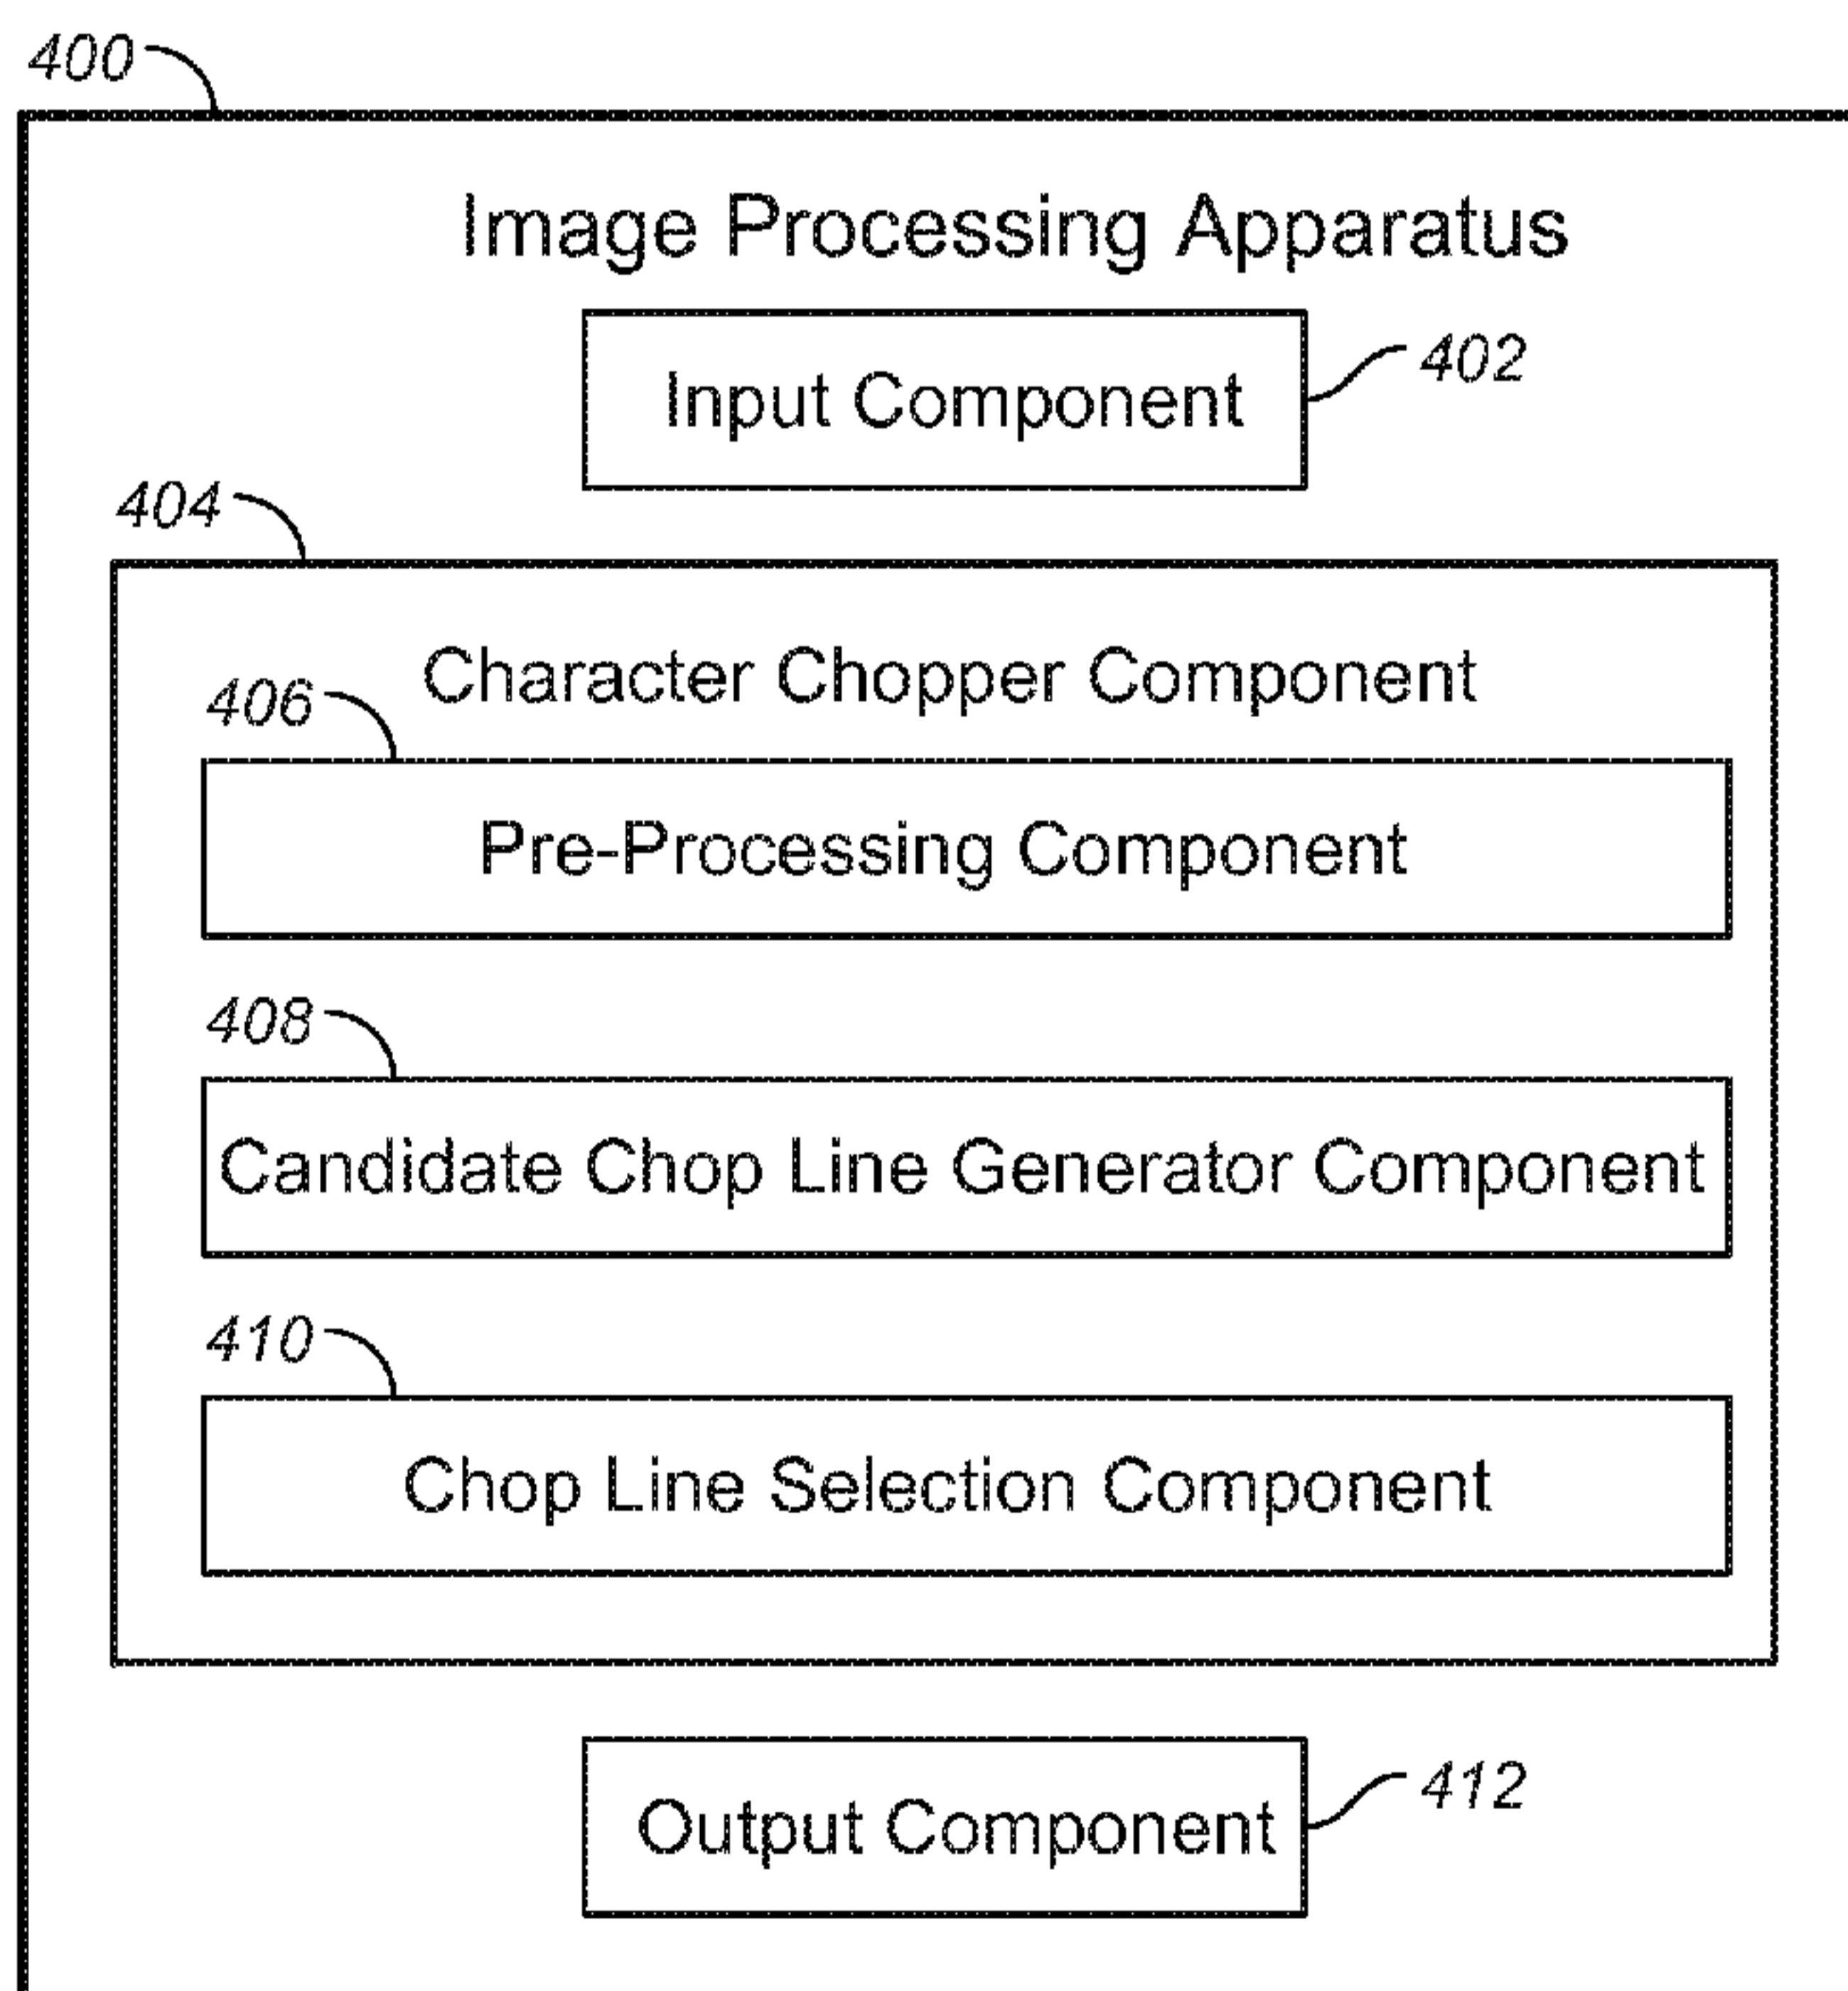


FIG. 4

(57) Abstract: An image processing apparatus is provided that includes a character chopper component that segments words into individual characters in a bitmap of a textual image undergoing an OCR process. The Character chopper component is configured to produce a set of (possibly curved) chop-lines which divide a bitmap of any given word into its individual character or glyph candidates. Cases where an input bitmap contains two separate words are handled by marking a place where those words should be split. The character segmentation algorithm computes the set of vertically oriented, curved chop-lines by considering glyph and background colors in a given word bitmap. The set is filtered afterwards using various heuristics, in order to preserve those lines that indeed do separate a word's glyphs and minimize the number of those that do not.

WO 2011/142977 A3



— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

— *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))*

Published:

— *with international search report (Art. 21(3))*

(88) Date of publication of the international search report:
12 January 2012

SEGMENTATION OF A WORD BITMAP INTO INDIVIDUAL CHARACTERS OR GLYPHS DURING AN OCR PROCESS

BACKGROUND

[0001] Optical character recognition (OCR) is a computer-based translation of an image of text into digital form as machine-editable text, generally in a standard encoding scheme. This process eliminates the need to manually type the document into the computer system. A number of different problems can arise due to poor image quality, imperfections caused by the scanning process, and the like. For example, a conventional OCR engine may be coupled to a flatbed scanner which scans a page of text. Because the page is placed flush against a scanning face of the scanner, an image generated by the scanner typically exhibits even contrast and illumination, reduced skew and distortion, and high resolution. Thus, the OCR engine can easily translate the text in the image into the machine-editable text. However, when the image is of a lesser quality with regard to contrast, illumination, skew, etc., performance of the OCR engine may be degraded and the processing time may be increased due to processing of all pixels in the image. This may be the case, for instance, when the image is obtained from a book or when it is generated by an imager-based scanner, because in these cases the text/picture is scanned from a distance, from varying orientations, and in varying illumination. Even if the performance of the scanning process is good, the performance of the OCR engine may be degraded when a relatively low quality page of text is being scanned.

[0002] One part of the OCR process identifies textual lines in a bitmap of a textual image. One component of the OCR engine segments each textual line with a series of chop lines that are located between adjacent characters or glyphs. Ideally, a single character or glyph is located between pair of adjacent chop lines. In many cases, however, it is difficult to segment words into individual symbols due to poor image quality, font weight, italic text, character shape, etc.

SUMMARY

[0003] An image processing apparatus is provided that includes a Character Chopper component that segments words into individual characters in a bitmap of a textual image undergoing an OCR process. The Character Chopper component is configured to produce a set of (possibly curved) chop-lines which divide a bitmap of any given word into its individual character or glyph candidates. Cases where an input bitmap contains two separate words are handled by marking a place where those words should be split. The character segmentation algorithm computes the set of vertically oriented, curved chop-lines by

51331-1304

considering glyph and background colors in a given word bitmap. The set is filtered afterwards using various heuristics, in order to preserve those lines that indeed do separate a word's glyphs and minimize the number of those that do not.

[0004] In one implementation, the segmentation process begins when a bitmap of an image is received. The image includes at least one textual line that is identified by a base-line and a mean-line. Next, a plurality of candidate chop lines are generated which each maximize a fitness function that increases as a total path lightness of the respective candidate chop lines increases and decreases as an intersection number increases. The intersection number denotes a number of white-to-black and black to white transitions (weighted according to their respective intensities) that the respective candidate chop line crosses. The fitness function thus selects candidate chop lines which contain more light pixels and fewer intersections with ink. At least one chop line is then selected from among the candidate chop lines. The chop line separates a pair of adjacent characters or glyphs in the textual line.

[0005] In one particular implementation, each of the candidate chop lines are assigned to one of a plurality of groups. The candidate chop lines in each group are all within a threshold logical distance from one another. The logical distance between two candidate chop lines is proportional to the sum of the inverted pixel values lying between the two candidate chop lines. A candidate chop line is selected from each group. The candidate chop line that is selected has a maximum confidence measure for its group. Any of the selected candidate chop lines are eliminated if their confidence measure is below a threshold level, thereby providing a first set of candidate chop lines. Pairs of candidate chop lines in the first set candidate chop lines are compared to one another. A first member of the pair is eliminated if its confidence measure is less than a confidence measure of a second member of the pair and if a product of the confidence measure of the second member and a decreasing function of the logical distance between the first and second members is greater than the confidence measure of the second member. In this way a second set of candidate chop lines is obtained which correspond to the final plurality of chop lines.

51331-1304

[0005a] According to one aspect of the present invention, there is provided an apparatus that generates characters or glyphs from a bitmap of text, comprising: a processor; memory, when operatively coupled to the processor, implementing: an input component for receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line; a character chopper component that includes a candidate chop line generator component for generating a plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line, wherein the candidate chop line generator component is configured to produce a candidate chop line through each pixel in at least one row extending along and within the textual line, said character chopper component further including a chop line selection component for selecting a subset of the candidate chop lines that correspond to the plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line; and an output component that applies the chop line to the textual line to produce the characters or glyphs, wherein the candidate chop line generator component generates candidate chop lines that each maximize a fitness function that increases as a total path lightness of the respective candidate chop lines increases and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses.

[0005b] According to another aspect of the present invention, there is provided a method for segmenting words in a bitmap of an image into characters or glyphs, comprising: receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line; generating a plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line by producing a candidate chop line through each pixel in at least one row extending along and within the textual line; and selecting a subset of the candidate chop lines that correspond to the plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line; wherein generating the plurality of chop lines includes generating the candidate chop lines so that they each maximize a fitness function that increases as a total path lightness of the respective candidate chop lines increases and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses.

51331-1304

[0005c] According to still another aspect of the present invention, there is provided a method for segmenting words in a bitmap of an image into characters or glyphs, comprising: receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line; generating a plurality of candidate chop lines that each maximize a fitness function
5 that increases as a total path lightness of the respective candidate chop lines increase and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses; selecting at least one chop line from among the candidate chop lines, wherein the chop line separates a pair of adjacent characters or glyphs in the textual line; applying the
10 chop line to the textual line to produce the characters or glyphs from the received bitmap of the image.

[0005d] According to yet another aspect of the present invention, there is provided a non-transitory computer-readable medium, having stored thereon, computer executable instructions, that when executed, perform a method according to a method described above.

[0006] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

51331-1304

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 shows one illustrative example of a system for performing optical character recognition (OCR) of a textual image.

5 [0008] FIG. 2 shows a high level flowchart of the character chopping process for segmenting each textual line with a series of chop lines that are located between adjacent characters or glyphs.

[0009] FIG. 3a shows the base-line (bottom horizontal line) and the mean-line (upper horizontal line) of a textual line that includes the words "of ground."

10 [0010] FIG. 3b shows a series of chop lines that have been calculated for the textual line shown in FIG. 3a.

[0011] FIG. 4 shows one example of an image processing apparatus that segments words into individual characters in a bitmap of a textual image undergoing an OCR process.

[0012] FIG. 5 shows a sequence in which the optimal half-line is determined for the upper part of the bigram "fy."

15 [0013] FIG. 6 shows a series of twelve consecutive candidate chop lines through the word "Larry" for the purpose of illustrating the concept of logical distance.

[0014] FIG. 7 illustrates a pre-processing step in which adjacent characters in a bitmap of a word touch one another.

20 [0015] FIG. 8 shows one example of a method for segmenting words in a bitmap of an image into characters or glyphs.

DETAILED DESCRIPTION

[0016] FIG. 1 shows one illustrative example of a system 5 for performing optical character recognition (OCR) of a textual image. The system 5 includes a data capture arrangement (e.g., a scanner 10) that generates an image of a document 15. The scanner 10 may be an imager-based scanner which utilizes a charge-coupled device as an image sensor to generate the image. The scanner 10 processes the image to generate input data, and transmits the input data to a processing arrangement (e.g., an OCR engine 20) for character recognition within the image. In this particular example the OCR engine 20 is incorporated into the scanner 10. In other examples, however, the OCR engine 20 may be a separate unit such as stand-alone unit or a unit that is incorporated into another device such as a PC, server, or the like.

[0017] The OCR engine 20 receives a textual image as a bitmap of text lines. The image may be a scanned image of text or a digital document such as a PDF or Microsoft WordTM document, where input data is already available. If such input data is not available, the

paragraph detection component of an OCR engine is assumed to receive as input from other components of the OCR engine an image in which a list of textual lines (and possibly words) have already been identified for each page of the image. Individual lines and words are defined by line and word bounding boxes, respectively. The bounding boxes themselves are defined by a set of coordinates that are established for each.

[0018] As previously mentioned, one component of the OCR engine 20 segments each textual line with a series of chop lines that are located between adjacent characters, glyphs or ligatures, which for the purposes herein will be treated as interchangeable with one another. For convenience, the term “glyph” may from time to time be used herein to mean a glyph, a character and/or a ligature. Ideally, a single glyph is located between pair of adjacent chop lines.

[0019] In many cases, however, it is difficult to segment words into individual characters, glyphs or ligatures due to poor image quality, font weight, italic text, character shape, etc. What makes the word segmentation even more difficult is that sometimes there is no clear spatial distinction between individual glyphs in a given word’s bitmap. For instance, characters can be joined, or an entire word can be underlined.

[0020] As detailed below, a character chopper component is configured to produce a set of (possibly curved) chop-lines which divide a bitmap of any given word into its individual character or glyph candidates. Cases where an input bitmap contains two separate words are handled by marking a place where those words should be split. The character segmentation algorithm computes the set of vertically oriented, curved chop-lines by considering glyph and background colors in a given word bitmap. The set is filtered afterwards using various heuristics, in order to preserve those lines that indeed do separate a word’s glyphs and minimize the number of those that do not.

[0021] FIG. 2 shows a high level flowchart of the character chopping process for segmenting each textual line with a series of chop lines that are located between adjacent characters or glyphs. The process begins at block 210 when a bitmap of an image is received. The bitmap includes at least one textual line of words. The textual line is identified by a base-line and a mean-line.

[0022] The base line is defined as a horizontal line passing through the bottom ends of a majority of the characters in a line of text (excluding descenders). Likewise, the “mean-line” is defined as a horizontal line which passes through the top ends of a majority of the characters in a line of text (excluding ascenders). The “x-height” is defined as the vertical distance between the base-line and the mean-line, which corresponds to the height of a

majority of lowercase letters in the line (excluding non-descenders and non-ascenders). FIG. 3a shows the base-line (bottom horizontal line) and the mean-line (upper horizontal line) of a textual line that includes the words “of ground.”

[0023] After the bitmap is received, the process continues at block 220 where candidate chop lines are generated. The candidate chop lines each separate a pair of adjacent characters or glyphs in the textual line. A confidence measure or level is calculated for each candidate chop line which reflects how accurately the chop line separates a character or glyph. FIG. 3b shows a series of chop lines that have been calculated for the textual line shown in FIG. 3a. In FIG. 3b, the darker the line the higher the confidence level. FIG. 3b also shows an inter-word chop line between the words “of” and “ground.” A subset of the candidate chop lines are selected which are most likely to correspond to actual inter-character or inter-glyph chop lines.

[0024] Finally, at block 230 the process outputs the subset of chop lines for each textual line. The output includes the coordinates of each textual line, the candidate chop lines and their confidence measures. This output may be used to apply the chop line to the textual line to identify or otherwise produce the characters or glyphs.

[0025] FIG. 4 shows one example of an image processing apparatus 400 that segments words into individual characters in a bitmap of a textual image undergoing an OCR process. The apparatus 400, which may be incorporated in an OCR engine, includes an input component 402 for receiving an input image having textual lines which are identified by a base-line and a mean-line that have been calculated in a prior component of the OCR engine. The apparatus also includes a character chopper component 404. The character chopper component 404 includes one or more pre-processing components 406 and a candidate chop line generator component 408 for generating a plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line. The character chopper component 404 also includes a chop line selection component 410 for selecting a subset of the candidate chop lines that correspond to the plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line. The apparatus 400 also includes an output component 412 that generates the information concerning the chop lines which have been generated in a form that allows it to be employed by subsequent components of the OCR engine. The subsequent component may be, for instance, a word search component that employs an algorithm such as a beam search algorithm that recognizes the words that the characters make up.

[0026] As mentioned above, the output from the system is a reasonably small set of (possibly curved) chop-lines extending from a word's top down to its bottom. While this set should properly chop the word's bitmap into individual character or glyph bitmaps, it is possible (although not desirable) for the set to have some additional (i.e., false) chop-lines (i.e. chop lines that chop some individual glyphs). If the system is "sure" that some particular chop-line in the output set is indeed an inter-glyph chop, it will flag it in an appropriate manner. Sometimes, however, due to word-segmentation errors, it may be possible that two consecutive words in a textual line received as part of the input bitmap to the system are treated as one bitmap. In such a case, the output set of chop lines will contain a line that splits the leftmost glyph of the first word and the rightmost glyph of the second word. The system will generally flag such a chop line as a potential inter-word chop.

[0027] An overview of the process by which chop lines are determined will now be presented. After this overview a more detailed discussion will be presented of a particular algorithm that may be employed to implement the process. This presentation will include a number of examples to further facilitate an understanding of the process.

Chop Line Computation Overview

[0028] Each chop-line determined by the character chopper component can be constructed using a number of anchor points, which will be defined below. While the number can vary from 0 to h , where h is bitmap height of the input image in pixels, the following examples will employ only a single anchor point. Given an input bitmap of height h and width w , the character chopper component computes a set of candidate chop-lines in two steps.

[0029] In the first step, an initial set of candidate chop-lines is computed using dynamic programming. One way to accomplish this is as follows. For each pixel along a central row of a word in a textual line the character chopper computes the most optimal vertically oriented chop-line containing that pixel. The central row is a row of pixels lying exactly between the base-line and the mean-line. Each of these pixels along the central row serves as the anchor point for its respective chop line. The chop line is constructed in such a way that it selects one and only one pixel from each row of the bitmap above and below the pixel on the central row. Moreover, the x-coordinate values of two pixels in a chop line selected from adjacent rows cannot differ by more than 1 pixel. This ensures that the chop line is a single, fully connected line. Since a chop line is computed for each pixel along the central row, w chop lines will be generated for an input bitmap having a width of w pixels. These chop lines constitute the set of candidate chop lines. Each candidate chop line will have a height of h pixels and will therefore contain exactly h pixels.

[0030] Once the w candidate chop lines have been computed, the second step of the chop line computation process is performed. In this step at least one subset of chop lines is selected from among the total set of candidate chop lines such that, for each pair of adjacent characters or glyphs in the word's bitmap, there exists at least one chop line in the subset that properly separates them. The process of selecting the proper subset of candidate chop lines which separate characters or glyphs can be computationally difficult. Accordingly, a chop line filtering process is employed in which various heuristics are used to reduce the number of candidate chop lines. One example of this filtering process will also be described in more detail below.

10 Computation of Candidate Chop Lines

[0031] As mentioned above, given an input bitmap h pixels in height and w pixels in width, a set of w candidate chop-lines is computed. The set is bijectively mapped to central row pixels – j^{th} line ($0 \leq j < w$) corresponds to j^{th} pixel of the central row. The line computation may be accomplished using dynamic programming.

15 [0032] Each candidate chop line that is computed is optimized to maximize the likelihood that it represents an actual chop line between characters or glyphs. That is, out of all the potential lines that may be formed through an anchor pixel, the candidate chop line represents the line that is most likely to represent an actual chop line. In order to identify such optimal chop lines, a line “optimality” criterion is chosen.

20 [0033] In one implementation the optimality criterion is selected to be a function of two properties of each line: its “path lightness,” which represents an overall luminance of the pixels that make up the line, and its “intersection count,” which reflects the number of white-to-black and black-to-white transitions that the line exhibits. This number may be weighted to take into account the intensity of the transition so that more intense transitions are given more weight. In other words, the optimal line through an anchor point that is selected as the candidate chop line is that line which contains more light pixels and fewer intersections with ink, which defines the foreground text in an image. A brief explanation of the rationale for using both components of this optimality criterion will now be presented.

[0034] A criterion that solely examines path lightness is not a good optimality criterion
30 because in some cases (such as when two adjacent glyphs touch one another) an intersection cannot be avoided. For instance, in one case a line may cut the letter ‘o’ in half, thus having two intersections with a single glyph while crossing a relatively small number of black pixels (due to the printing quality of the letter ‘o’). In another case a line may properly chop the “ry” string, which in some fonts is printed as a single glyph (with ‘r’ and ‘y’ joined at

the top), and thus the line only intersects once with foreground ink. However, the line in the second case may potentially cross the same number of black pixels as the first line. Clearly, in order to determine that the line in the second case has a higher probability of being an actual chop-line, path lightness cannot be the only criterion.

5 [0035] Similarly, a criterion that only examines the number of intersections is also not a good optimality criterion. For instance, in one case a chop-line may enter the letter “l” at its top, and exit it at the bottom, making one white-to-black and one black-to-white intersection. In another case a chop line may intersect two underlined glyphs that touch one another where they are underlined. The lines in the first and second case will thus have the
10 same intersection count. Clearly, in order to determine that the line in the second case has a higher probability of being an actual chop-line, intersection count cannot be the only criterion.

[0036] A candidate chop-line’s path lightness may be computed by summing weighted pixel values across the line. The weighting factors arise from differences in background
15 color and are produced by a background filtering component in the OCR engine. The intersection count can be computed by summing weighted luminance differences of spatially close pixels on a chop-line. An illustrative method for determining the path lightness and intersection count and for combining them into the optimality criterion will be provided below.

20 [0037] In one implementation two halves of a candidate chop line are computed separately from one another and then combined. One half may correspond to the portion of the chop line above the central row and the other half may correspond to the portion of the chop line below the central row. In particular, a chop-line corresponding to pixel p_i on the central row is composed of an upper and lower half line. The upper half-line starts at p_i and ends at
25 some (determined by the half-line’s shape) pixel at the bitmap’s first or top row. The lower half-line starts at p_i as well, and ends at some pixel at the bitmap’s last or bottom row. The two half-lines are computed separately using the optimality criterion described in the following section and joined afterward into a final candidate chop-line.

[0038] The upper and lower half-lines are both computed in a similar fashion. Thus, only
30 the computation of the upper half-line will be described herein. The half-line is computed row-after-row (starting from the first or top row) using dynamic programming. To describe the half-line computation, the following terms/data structures are defined (consider $0 \leq i \leq \frac{h}{2}, 0 \leq j < w$):

- *Optimal line (i, j)*: A line computed by the algorithm, having its first pixel somewhere on the bitmap's first or top row, and its last pixel in the bitmap's i^{th} row and j^{th} column.
- *Optimal line (i, j, k)*: Value of k^{th} pixel in *Optimal line (i, j)*
- 5 • $intersectionCount[i, j] = \sum_{k=1}^j abs(Optimal\ line\ (i, j, k) - Optimal\ line\ (i, j, k - 1))$
- $lightness[i, j] = \sum_{k=0}^j Optimal\ line\ (i, j, k)$
- $FitnessFunc(a, b, B_{CLAMP}) = a * \left(1.0 - \frac{\min(b, B_{CLAMP})}{B_{CLAMP}}\right) \dots (1)$

where B_{CLAMP} is a constant for normalization purposes

- 10 • $fitness[i, j]$: Measure of *optimal line (i, j)*'s quality. It is computed from the following formula:

$$fitness[i, j] = FitnessFunc(lightness[i, j], intersectionCount[i, j], INTERSECTIONMAX) \dots (2)$$

[0039] Clearly, the higher the lightness, the higher the fitness. On the other hand, the higher the (normalized) intersection count, the lower the fitness becomes.

The INTERSECTIONMAX clamp constant B_{CLAMP} is chosen empirically.

- 15 [0040] Once the fitness is computed for the i^{th} row, the computation for the $(i+1)^{\text{th}}$ row is done using the following formula:

$$fitness[i + 1, j] = \max \left\{ \begin{array}{l} FitnessFunc(lightness[i, j - 1] + img[i + 1, j], intersectionCount[i, j - 1] + abs(img[i + 1, j] - img[i, j - 1]), M) \\ FitnessFunc(lightness[i, j] + img[i + 1, j], intersectionCount[i, j] + abs(img[i + 1, j] - img[i, j]), M) \\ FitnessFunc(lightness[i, j + 1] + img[i + 1, j], intersectionCount[i, j + 1] + abs(img[i + 1, j] - img[i, j + 1]), M) \end{array} \right\} \dots (3)$$

Where:

- img is the input bitmap
 - M is an empirically chosen clamp constant (INTERSECTIONMAX in (2))
- 20 [0041] The reasoning behind this is rather simple: for each pixel in the row currently being processed, the algorithm analyzes whether it is "cheaper" (i.e., maximizes the fitness function) to reach that pixel from its upper-left, upper, or upper-right neighbor. After processing the central row, the $fitness[h/2, j]$ ($0 \leq j < w$) will reflect the line quality measurement values for all of the w upper half-lines, and each one of them will be
- 25 "optimal" starting from its own (and unique) pixel on the central row.

[0042] There may be an additional penalty for diagonal (upper-left and upper-right) transitions in eq. 3, which is omitted for clarity. The penalty coefficients may be different if

the analyzed word is in italic style. The detection of italic text will be described in a separate section below.

When finding the value of the $fitness[i + 1, j]$ in formula (3), the $lightness[i + 1, j]$ and $intersectionCount[i + 1, j]$ are computed depending on which term from eq. (3) is chosen
5 as the maximum one.

[0043] The algorithm for determining a half-line is illustrated in connection with FIG. 5, which shows an original image (the upper part of the bigram “fy”) in the top left portion of the figure. A single target pixel 200 is selected (top middle portion of FIG. 5). As shown in the top right portion of FIG. 5, the optimal path through the pixel 200 will contain either the
10 neighboring pixel to its top-left (pixel 202), top (pixel 204) or top-right (pixel 206). Since the chop line computation process begins on the image’s top row and continues downward, it can be assumed that optimal chop line paths containing those three pixels have already computed (line 220 in bottom-left, line 230 in bottom-middle and line 240 in bottom-right pictures in FIG. 5). For those three paths that have already been calculated, plus the target
15 pixel, three 3 fitness functions are computed and compared to one another. The path with the largest fitness is selected as the optimal path for the target pixel. This process continues for each pixel on the current row, which is then repeated until the last row is reached. In the half line approach the last row will be the central row. The entire process is then performed for the lower half of the input image (not shown in FIG. 5). Finally, for each pixel along the
20 central row, the upper and lower half-lines that have been calculated are joined together to form a candidate chop line.

[0044] After the candidate chop lines have been computed their respective confidences may be determined. One way to derive the confidence measure is as follows. First, the path
lightness is transformed so that it lies within some predefined range, which in this example
25 is chosen to be from 0.0 to 1.0. In theory, the value of the lightness cannot exceed the product of the maximum pixel value and the bitmap’s height, so this constant ($maxlight$) can be used as a scaling factor. Moreover, it is possible to empirically choose some non-zero value ($minlight$) such that lines with a lightness lower than this value can never be
30 considered as proper chop line candidates. For example, a line entering the “l” letter at its top, and exiting at its bottom, has, say, around 80% of its length in the dark area occupied by ink, and thus can be safely omitted from analysis. To put more emphasis on a value reflecting the line’s quality, lightness values between $minlight$ and $maxlight$ may be linearly interpolated to the (0.0, 1.0) range, while those below $minlight$ are clamped to zero. This

scaled lightness value is then multiplied by the term in parenthesis from eq. (1) describing the intersection count, which will be in the range of (0.0 to 1.0).

Candidate Chop Line Filtering

5 [0045] As previously mentioned, after the set of candidate chop lines has been determined, at least one subset of the lines is selected such that for each pair of adjacent characters or glyphs in a word's bitmap, there exists one and only one chop line in the subset that properly separates them. Since this can be a computationally intensive task, the candidate chop lines are filtered to reduce the size of the set as much as possible in a computationally efficient way, while retaining the correct subset of chop lines.

10 [0046] The filtering algorithm employs a metric that will be referred to as the chop-line logical distance. The chop-line logical distance is defined as the sum of the inverted pixel values from the input bitmap which lie between two chop-lines. Thus, the darker these pixels become, the greater the chop line logical distance between those two lines becomes. Conversely, if only a pure white background lies between some two particular chop lines,
15 the logical distance between them will be equal to zero.

[0047] The initial set of w candidate chop lines are first arranged in a strictly ascending order by the location of their anchor pixel along the input bitmap's central row. The lines are then arranged into groups based on their logical distance. Any lines that are within a certain threshold logical distance from line 0 are added to the first group. The first line that
20 exceeds the threshold is added to a second group. This first line in the second group is then used as a new reference line so that any subsequent lines that are within a second threshold logical distance of it are added to the second group. This process is repeated until the last line is reached. The threshold values are chosen empirically. Before comparing logical distances to the threshold, they are normalized (divided by their squared x-height value), to
25 compensate for different input bitmap sizes.

[0048] Next, a line is selected from each group which has the greatest confidence measure. All the other lines are filtered-out. Any candidate chop line in this newly established set of lines is also filtered out if its confidence is below some threshold value. A second filtering step is then performed by examining the local neighborhood (in terms of logical distance)
30 surrounding each member of this newly formed group. This neighborhood is searched for any additional lines that can be filtered out. For instance, suppose two lines A and B which are both members of the newly formed group are in the same local neighborhood and $confidence_A > confidence_B$. Depending on the logical distance of A and B and their confidences ratio, line B may or may not be filtered out. In particular, $confidence_A$ is

multiplied by a decreasing (e.g., exponentially) function of the logical distance between lines A and B after the logical distance has been normalized within the range of 0.0 to 1.0. If this product is higher than the value of *confidence_B*, line B will be filtered out. The logic behind this approach is that “weak” lines lying close to “strong” lines should be excluded
5 from further analysis.

[0049] In some cases the filtering algorithm may include a weighting function to provide a logical distance weighting measure. The reason for this is that it may be useful to differentiate if ink lying between two chop-lines is situated near the input bitmap’s central row, or somewhere near its top or bottom. In particular, two candidate chop lines that have
10 ink between them located at the input bitmap’s top or bottom may be marked as logically “farther” than some other two lines having the same amount of ink or “blackness” between them, but which are located directly along the central row. This is an empirically derived heuristic that can be easily illustrated. For instance, two lines are not “logically close” if there is a punctuation glyph (thus, a small amount of “blackness”) between them. Since
15 punctuation glyphs lie either near the input bitmap’s top or near its bottom, this heuristic makes intuitive sense. The modification to the original logical distance computation algorithm that is needed to take such a weighting factor into account is straightforward – the computed sum of the pixel values for each row is simply multiplied by a specific, row-wise scaling factor.

[0050] The filtering algorithm may also reduce in weight or even filter out chop lines that are too curvy (where curviness is computed as an x-coordinate difference between the line’s top and bottom pixel). Conversely, the filtering algorithm may give lines that are “straight” additional weight.

[0051] It is sometimes possible that a certain number of low-confidence lines exist between
25 two high-confidence lines even after the filtering described above has been performed. This may arise, for example, when a collection of lines pass through the diagonal part of the letter “N.” To account for this situation, only a small, fixed number of low-confidence lines may be allowed to exist between two “reasonably close” high-confidence ones.

[0052] The concept of logical distance may be illustrated in connection with FIG. 6, which
30 shows a series of 12 consecutive candidate chop lines through the word “Larry.” That is, each candidate chop line has the x-coordinate of its anchor point increased by 1 with respect to the preceding candidate chop line. The sequence shown in FIG. 6 follows normal reading order beginning with the upper left word and continuing from left to right and top to bottom. It can be observed that there is a significant amount of ink between lines 1 and 2 (the lower

part of “a” letter), and between lines 11 and 12 (the lower part of “r” letter). However, there is not much ink between lines 2 and 3, lines 2 and 4, lines 2 and 5, and so on, all the way to lines 2 and 11. Accordingly, lines 2 and 11 are still logically close to one another. Lines 2 to 11 will therefore likely belong to the same group, and thus all the lines that are illustrated
 5 except the maximum-confidence one will be filtered out.

[0053] In addition to a confidence measure, additional information or annotations may be determined which are then associated with the candidate chop lines. For instance, a candidate chop line may be designated as an actual chop line if it doesn't overlap with the background, and if its nearest surrounding neighborhood from the un-filtered set of
 10 candidate chop lines consists only of high confidence lines that do not overlap with the background. The size of the nearest surrounding neighborhood may be chosen empirically, as may high confidence threshold needed to designate a line as a high confidence line. As another example, a candidate chop line may be designated as an inter-word chop line if it meets the same conditions required for an actual chop line, except that the size of its nearest
 15 surrounding neighborhood exceeds the size of the surrounding neighborhood for a line designated as an actual chop line by a certain amount (which may also be empirically determined).

Pre-processing Steps

[0054] A number of pre-processing steps may be performed before segmenting words in a
 20 textual line into individual characters or glyphs. For example, the chop lines for italic text will clearly have a different dominant slope from un-italicized text and thus text that is italicized may be recognized as such in a pre-processing step. Likewise, in some cases two characters or glyphs will touch one another, making the segmentation process more difficult. Thus it would be helpful to identify such characters or glyphs before determining
 25 the candidate chop lines during the segmentation process.

[0055] Regarding italicized text, the dominant slope of the chop-lines for italic text will clearly be different in comparison to normal text. The dominant chop-line slope may be favored by expanding eq. 3 to applying different penalty coefficients to chop lines that slope in different directions. However, before changing these coefficients, it is necessary to detect
 30 if a given word is in italic style.

[0056] The detection of italicized text can be performed as follows. Let $v_{proj}[j]$ be the sum of pixel values in the j^{th} column of a given word's bitmap. The projection will have a certain number of peaks and valleys arising from inter-glyph spacing and dominant vertical glyph parts respectively. Peaks and valleys are going to appear in alternating order. In the case of

“normal” text style, these “oscillations” are generally more visible than in italic text because the dominant parts of a character or glyph and the spaces between them are not as vertical as in the case of normal, un-italicized text.

[0057] As an example, if the same word in one style is skewed to the left by a certain angle with respect to the same word in another style, two cases may arise:

1. The original word is printed in normal style. The skewed version of the word will look similar to italic style, and the *vproj* oscillations in the first case will be more visible.
2. The original word is printed in italic style. The skewed version of the word will look similar to normal style, and the *vproj* oscillations in the second case will be more visible.

[0058] These heuristics may be used to detect italic style. There are many ways to compute a value that reflects oscillation strength. One of them may be:

$$osc_strength = \sum_j (vproj[j] - vproj_avg)^2$$

Where *vproj_avg* is an average value of the *vproj* array.

[0059] In regard to the case of two characters or glyphs that touch one another, the segmentation algorithm described above will fail if two characters or glyphs share a large vertical portion of an ink. For example, a pair of “o”s can sometimes appear “glued” to one another. To handle those cases, a simple pre-processing stage is added before chop-line computation. The stage operates in the following way:

1. Top and bottom text contours are computed from the input bitmap. Let the bitmap’s width be denoted as *w*. For each *x* coordinate ($x = 0..w-1$), the top contour tracks the *y*-coordinate of the first non-background pixel in a row *x*, looking from the bitmap’s top. Likewise, the bottom contour tracks the *y*-coordinate of the first non-background pixel in a row *x*, looking from the bitmap’s bottom.
2. Due to the shape of the character or glyph, there will be a local minima in the top contour (“valleys”) and a local maxima in the bottom contour (“hills”). If there is an *x* coordinate where a valley in the top contour is above the hill in the bottom contour, there is a high probability that this *x* coordinate corresponds to a place where two characters or glyphs touch one another.
3. Once these local minima and maxima are identified the shortest path between them is computed. Non-background pixels on the shortest path are then lightened in order to facilitate the identification of a potential chop-line candidate at that location. This modification to the input bitmap allows the chop-line computation stage to identify potential

chop-lines even in cases where two touching characters or glyphs would otherwise be difficult to segment.

[0060] The process of modifying the input bitmap in this manner is illustrated in FIG. 7. In FIG. 7 the original bitmap is shown on top. The bottom bitmap has been modified to lighten the non-background pixels along the line connecting the local minima and maxima in the top and bottom contours that arise between the bigrams “sc” “ss” and “se.”

[0061] FIG. 8 shows one example of a method 300 for segmenting words in a bitmap of an image into characters or glyphs. The method 300 begins in step 310 when a bitmap of an image is received. The image includes at least one textual line that is identified by a baseline and a mean-line. Next, in step 320 a plurality of candidate chop lines are generated which each maximize a fitness function that increases as a total path lightness of the respective candidate chop lines increases and decreases as an intersection number increases. The intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses. At least one chop line is selected from among the candidate chop lines in step 330. The chop line separates a pair of adjacent characters or glyphs in the textual line.

[0062] As used in this application, the terms "component," "module," "engine," "system," "apparatus," "interface," or the like are generally intended to refer to a computer-related entity, either hardware, a combination of hardware and software, software, or software in execution. For example, a component may be, but is not limited to being, a process running on a processor, a processor, an object, an executable, a thread of execution, a program, and/or a computer. By way of illustration, both an application running on a controller and the controller can be a component. One or more components may reside within a process and/or thread of execution and a component may be localized on one computer and/or distributed between two or more computers.

[0063] Furthermore, the claimed subject matter may be implemented as a method, apparatus, or article of manufacture using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof to control a computer to implement the disclosed subject matter. The term "article of manufacture" as used herein is intended to encompass a computer program accessible from any computer-readable device, carrier, or media. For example, computer readable media can include but are not limited to magnetic storage devices (e.g., hard disk, floppy disk, magnetic strips . . .), optical disks (e.g., compact disk (CD), digital versatile disk (DVD) . . .), smart cards, and flash memory devices (e.g., card, stick, key drive . . .). Of course, those skilled in the art

will recognize many modifications may be made to this configuration without departing from the scope or spirit of the claimed subject matter.

[0064] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in
5 the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

51331-1304

CLAIMS:

1. An apparatus that generates characters or glyphs from a bitmap of text, comprising:

a processor;

5 memory, when operatively coupled to the processor, implementing:

an input component for receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line;

a character chopper component that includes a candidate chop line generator component for generating a plurality of chop lines that each separate a pair of adjacent
10 characters or glyphs in the textual line, wherein the candidate chop line generator component is configured to produce a candidate chop line through each pixel in at least one row extending along and within the textual line, said character chopper component further including a chop line selection component for selecting a subset of the candidate chop lines that correspond to the plurality of chop lines that each separate a pair of adjacent characters or
15 glyphs in the textual line; and

an output component that applies the chop line to the textual line to produce the characters or glyphs,

wherein the candidate chop line generator component generates candidate chop lines that each maximize a fitness function that increases as a total path lightness of the
20 respective candidate chop lines increases and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses.

2. The apparatus of claim 1 wherein the candidate chop line generator component extends each of the character chop lines for additional pixel rows extending along and within
25 the textual line while continuing to maximize each character chop line.

51331-1304

3. The apparatus of claim 2 wherein the additional pixel rows include all pixel rows within the textual line.

4. The apparatus of claim 1 wherein the at least one row is a central row halfway between a base line and a mean line of the textual line and the character chopper component is
5 configured to separately produce candidate chop half-lines above and below the central row which are subsequently joined together.

5. The apparatus of claim 1 wherein the character chop line generator component generates a confidence measure associated with each chop line which reflects how accurately the chop line separates a character or glyph.

10 6. The apparatus of claim 5 wherein the confidence measure depends on a product of a scaled path lightness and an intersection number, wherein the scaled path lightness is assigned a value of zero for path lightness values below a threshold level.

7. The apparatus of claim 1 wherein the chop line selection component selects the subset of candidate chop lines by:

15 assigning each of the candidate chop lines to a plurality of groups, wherein the candidate chop lines in each group are all within a threshold logical distance from one another, said logical distance between two candidate chop lines being proportional to a sum of inverted pixel values lying between the two candidate chop lines;

20 selecting a candidate chop line from each group that has a maximum confidence measure;

eliminating any of the selected candidate chop lines if their confidence measure is below a threshold level to thereby provide a first set of candidate chop lines; and

25 comparing pairs of candidate chop lines in the first set candidate chop lines and eliminating a first member of a compared pair if its confidence measure is less than a confidence measure of a second member of the pair and if a product of the confidence measure of the second member and a decreasing function of the logical distance between the

51331-1304

first and second members is greater than the confidence measure of the second member to thereby provide a second set of candidate chop lines which are designated to be the subset of the candidate chop lines that correspond to the plurality of chop lines.

8. The apparatus of claim 7 wherein the logical distance between two candidate chop lines is proportional to a weighted sum of the inverted pixel values lying between the two candidate chop lines, wherein inverted pixel values of pixels closer to a central line through the textual line are weighted more than inverted pixel values of pixels more remote from the central line toward either the mean-line or the base-line.
9. The apparatus of claim 7 wherein the chop line selection component eliminates selected candidate chop lines that do not overlap any background regions of the image that do not contain text and for which there are no additional candidate chop lines within a specified distance of the selected candidate chop lines such that the additional candidate chop lines do not overlap any background regions of the image that do not contain text.
10. The apparatus of claim 1 wherein the character chopper component includes a first pre-processing component for detecting adjacent glyphs or characters that are in contact with one another such that they share a vertical portion of text, said detection component being configured to identify top and bottom contours defined by y-coordinates of non-background pixels in top and bottom rows of the image and detect a local minima in the top contour and a local maxima in the bottom contour having a common x-coordinate, said character detection component being further configured to lighten a value of pixels located on a vertical line connecting the local minima to the local maxima.
11. The apparatus of claim 10 wherein the character chopper component includes a second pre-processing component that detects italicized characters in the textual line based on a measure of oscillations in a sum of pixel values in columns of the image.
12. A method for segmenting words in a bitmap of an image into characters or glyphs, comprising:

51331-1304

receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line;

generating a plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line by producing a candidate chop line through each pixel
5 in at least one row extending along and within the textual line; and

selecting a subset of the candidate chop lines that correspond to the plurality of chop lines that each separate a pair of adjacent characters or glyphs in the textual line;

wherein generating the plurality of chop lines includes generating the candidate chop lines so that they each maximize a fitness function that increases as a total path lightness
10 of the respective candidate chop lines increases and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses.

13. The method of claim 12 further comprising extending each of the character chop lines for each pixel row extending along and within the textual line while continuing to
15 maximize each character chop line.

14. The method of claim 12 wherein the at least one row is a central row halfway between a base line and a mean line of the textual line and the character chopper component is configured to separately produce candidate chop half-lines above and below the central row which are subsequently joined together.

20 15. The method of claim 12 further comprising generating a confidence measure associated with each chop line which reflects how accurately the chop line separates a character or glyph, the confidence measure depending on a product of a scaled path lightness and an intersection number, wherein the scaled path lightness is assigned a value of zero for path lightness values below a threshold level.

25 16. A method for segmenting words in a bitmap of an image into characters or glyphs, comprising:

51331-1304

receiving a bitmap of an image comprising at least one textual line that is identified by a base-line and a mean-line;

generating a plurality of candidate chop lines that each maximize a fitness function that increases as a total path lightness of the respective candidate chop lines increase and decreases as an intersection number increases, wherein the intersection number denotes a number of white-to-black and black to white transitions that the respective candidate chop line crosses;

selecting at least one chop line from among the candidate chop lines, wherein the chop line separates a pair of adjacent characters or glyphs in the textual line;

applying the chop line to the textual line to produce the characters or glyphs from the received bitmap of the image.

17. The method of claim 16 wherein each of the candidate chop lines pass through a different pixel in at least one row extending along and within the textual line.

18. The method of claim 16 wherein the chop line selection is further arranged for:

assigning each of the candidate chop lines to a plurality of groups, wherein the candidate chop lines in each group are all within a threshold logical distance from one another, said logical distance between two candidate chop lines being proportional to a sum of inverted pixel values lying between the two candidate chop lines;

selecting a candidate chop line from each group that has a maximum confidence measure;

eliminating any of the selected candidate chop lines if their confidence measure is below a threshold level to thereby provide a first set of candidate chop lines; and

comparing pairs of candidate chop lines in the first set candidate chop lines and eliminating a first member of the pair if its confidence measure is less than a confidence measure of a second member of the pair and if a product of the confidence measure of the

51331-1304

second member and a decreasing function of the logical distance between the first and second members is greater than the confidence measure of the second member to thereby provide a second set of candidate chop lines which are designated to be the subset of the candidate chop lines that correspond to the at least one chop line.

- 5 19. A non-transitory computer-readable medium, having stored thereon, computer executable instructions, that when executed, perform a method according to any one of claims 12 to 18.

1/4


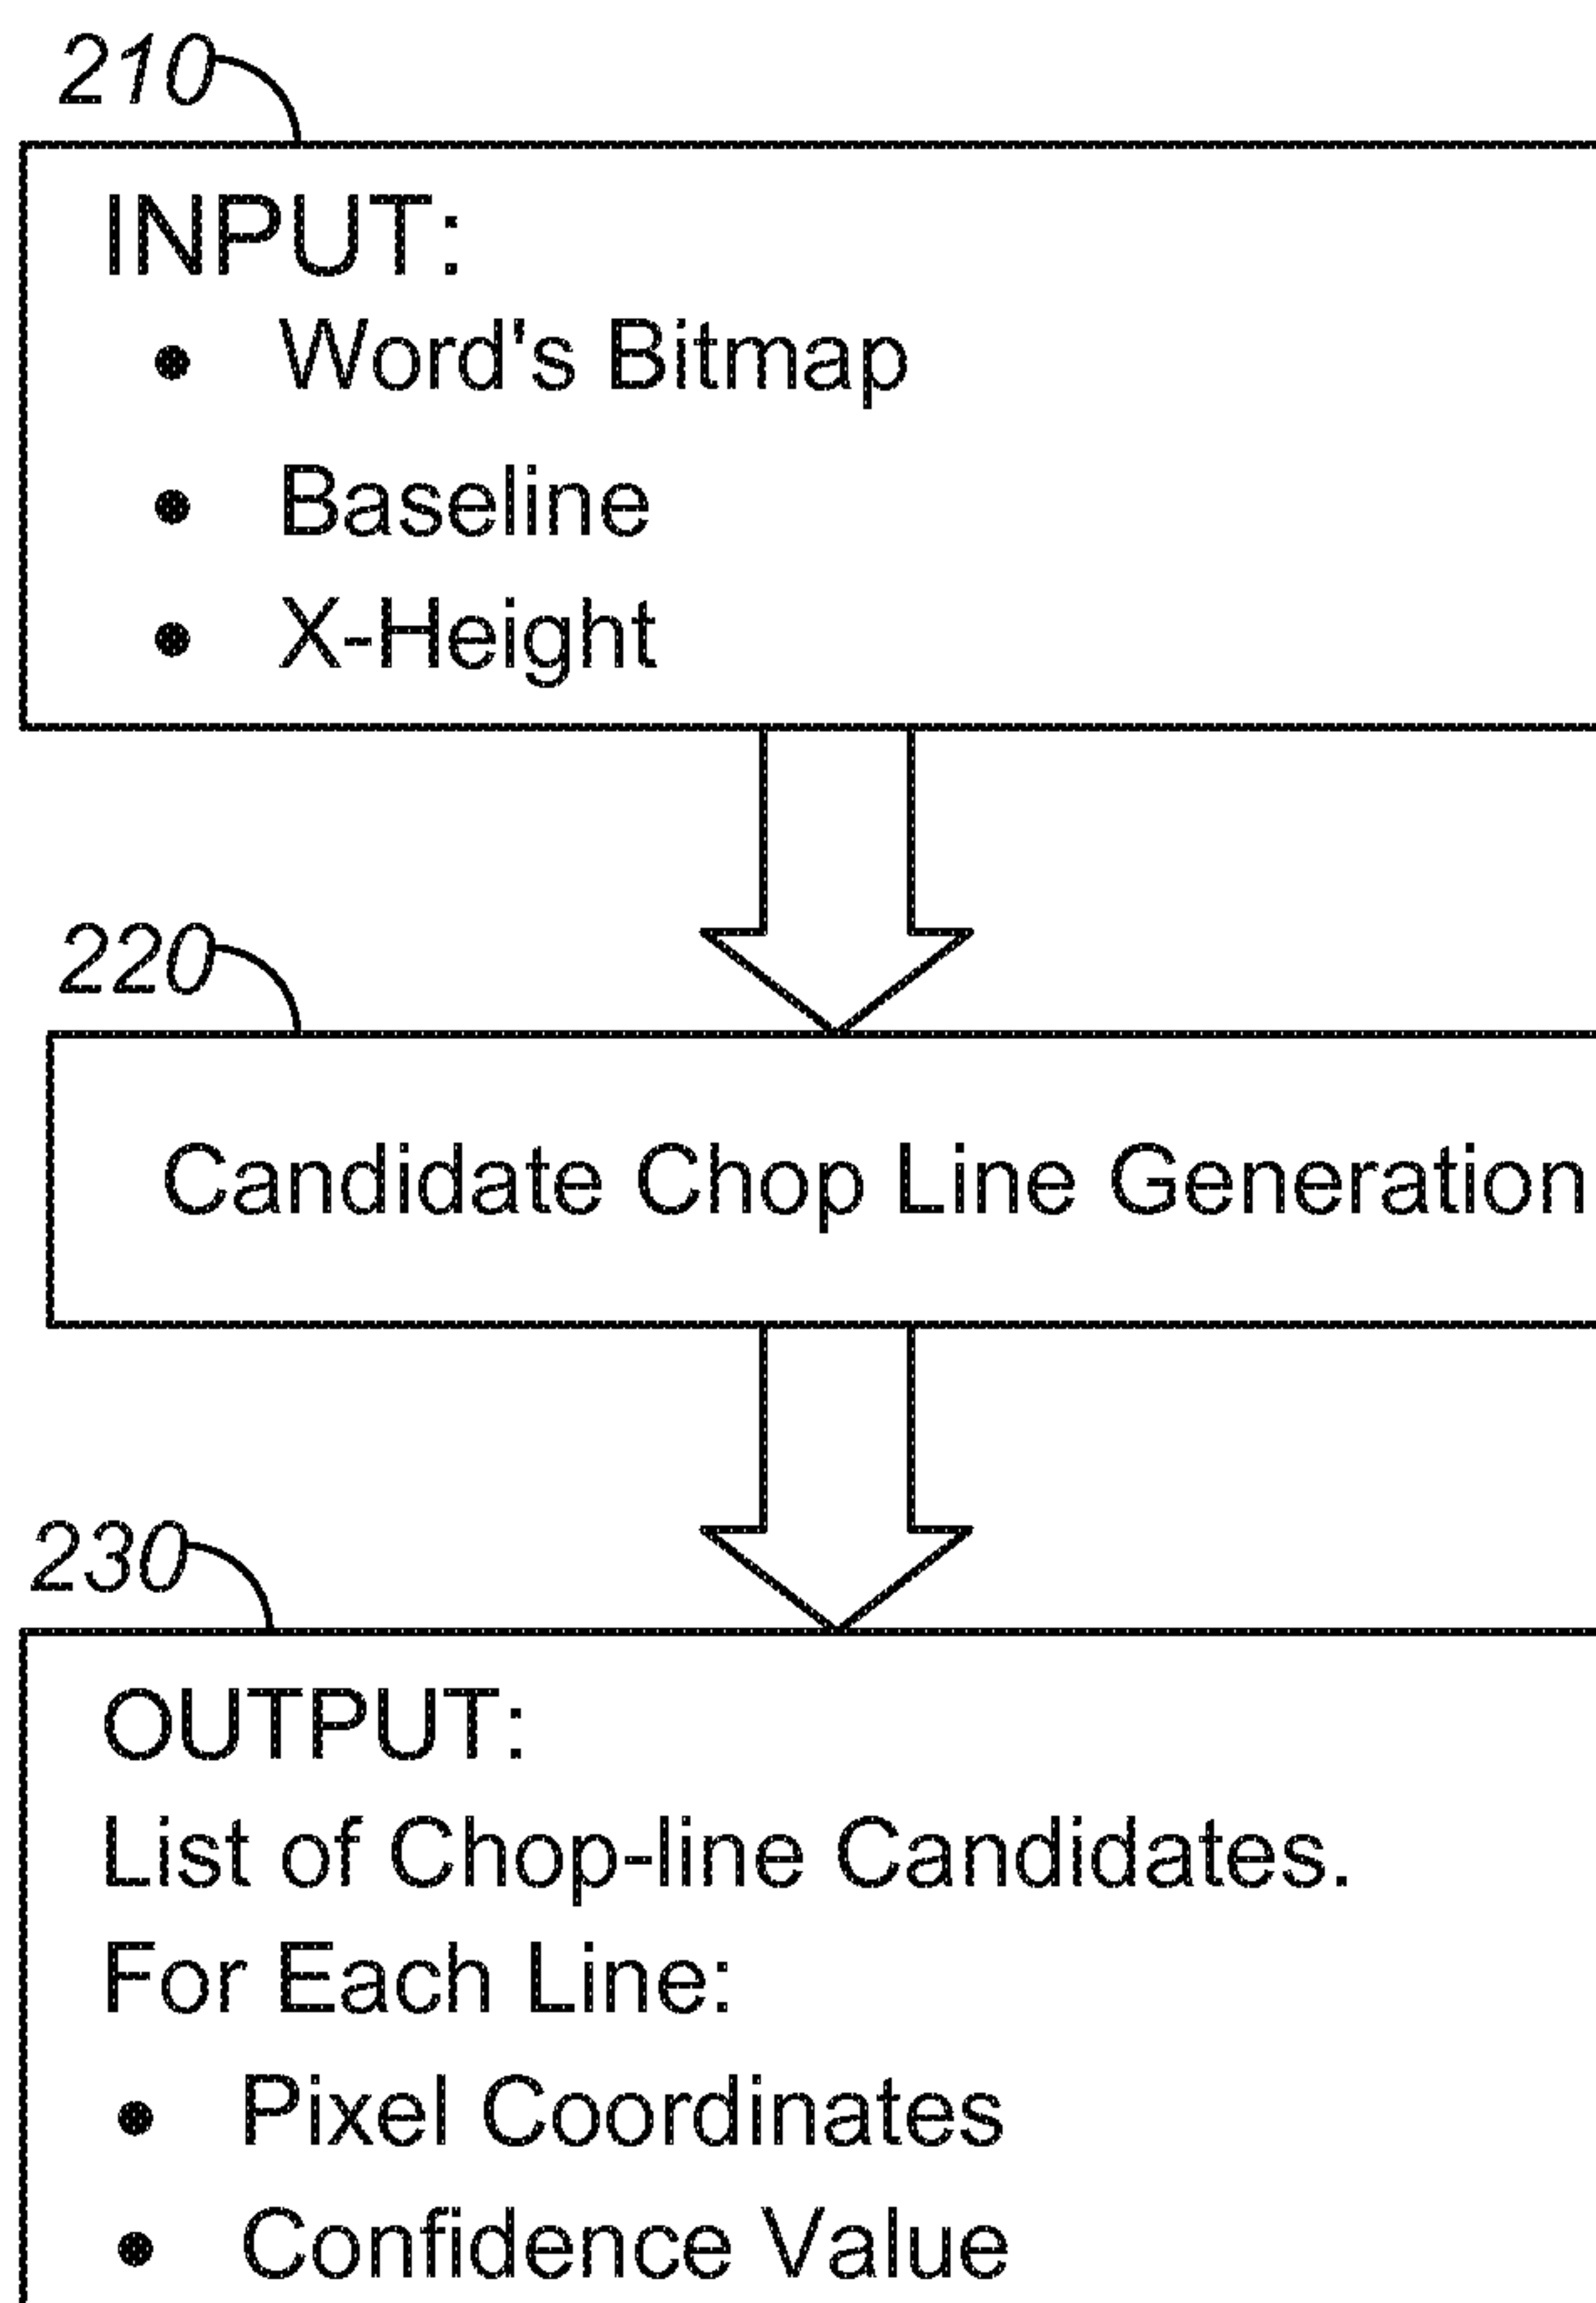
**FIG. 1****FIG. 2**



FIG. 3a




FIG. 3b

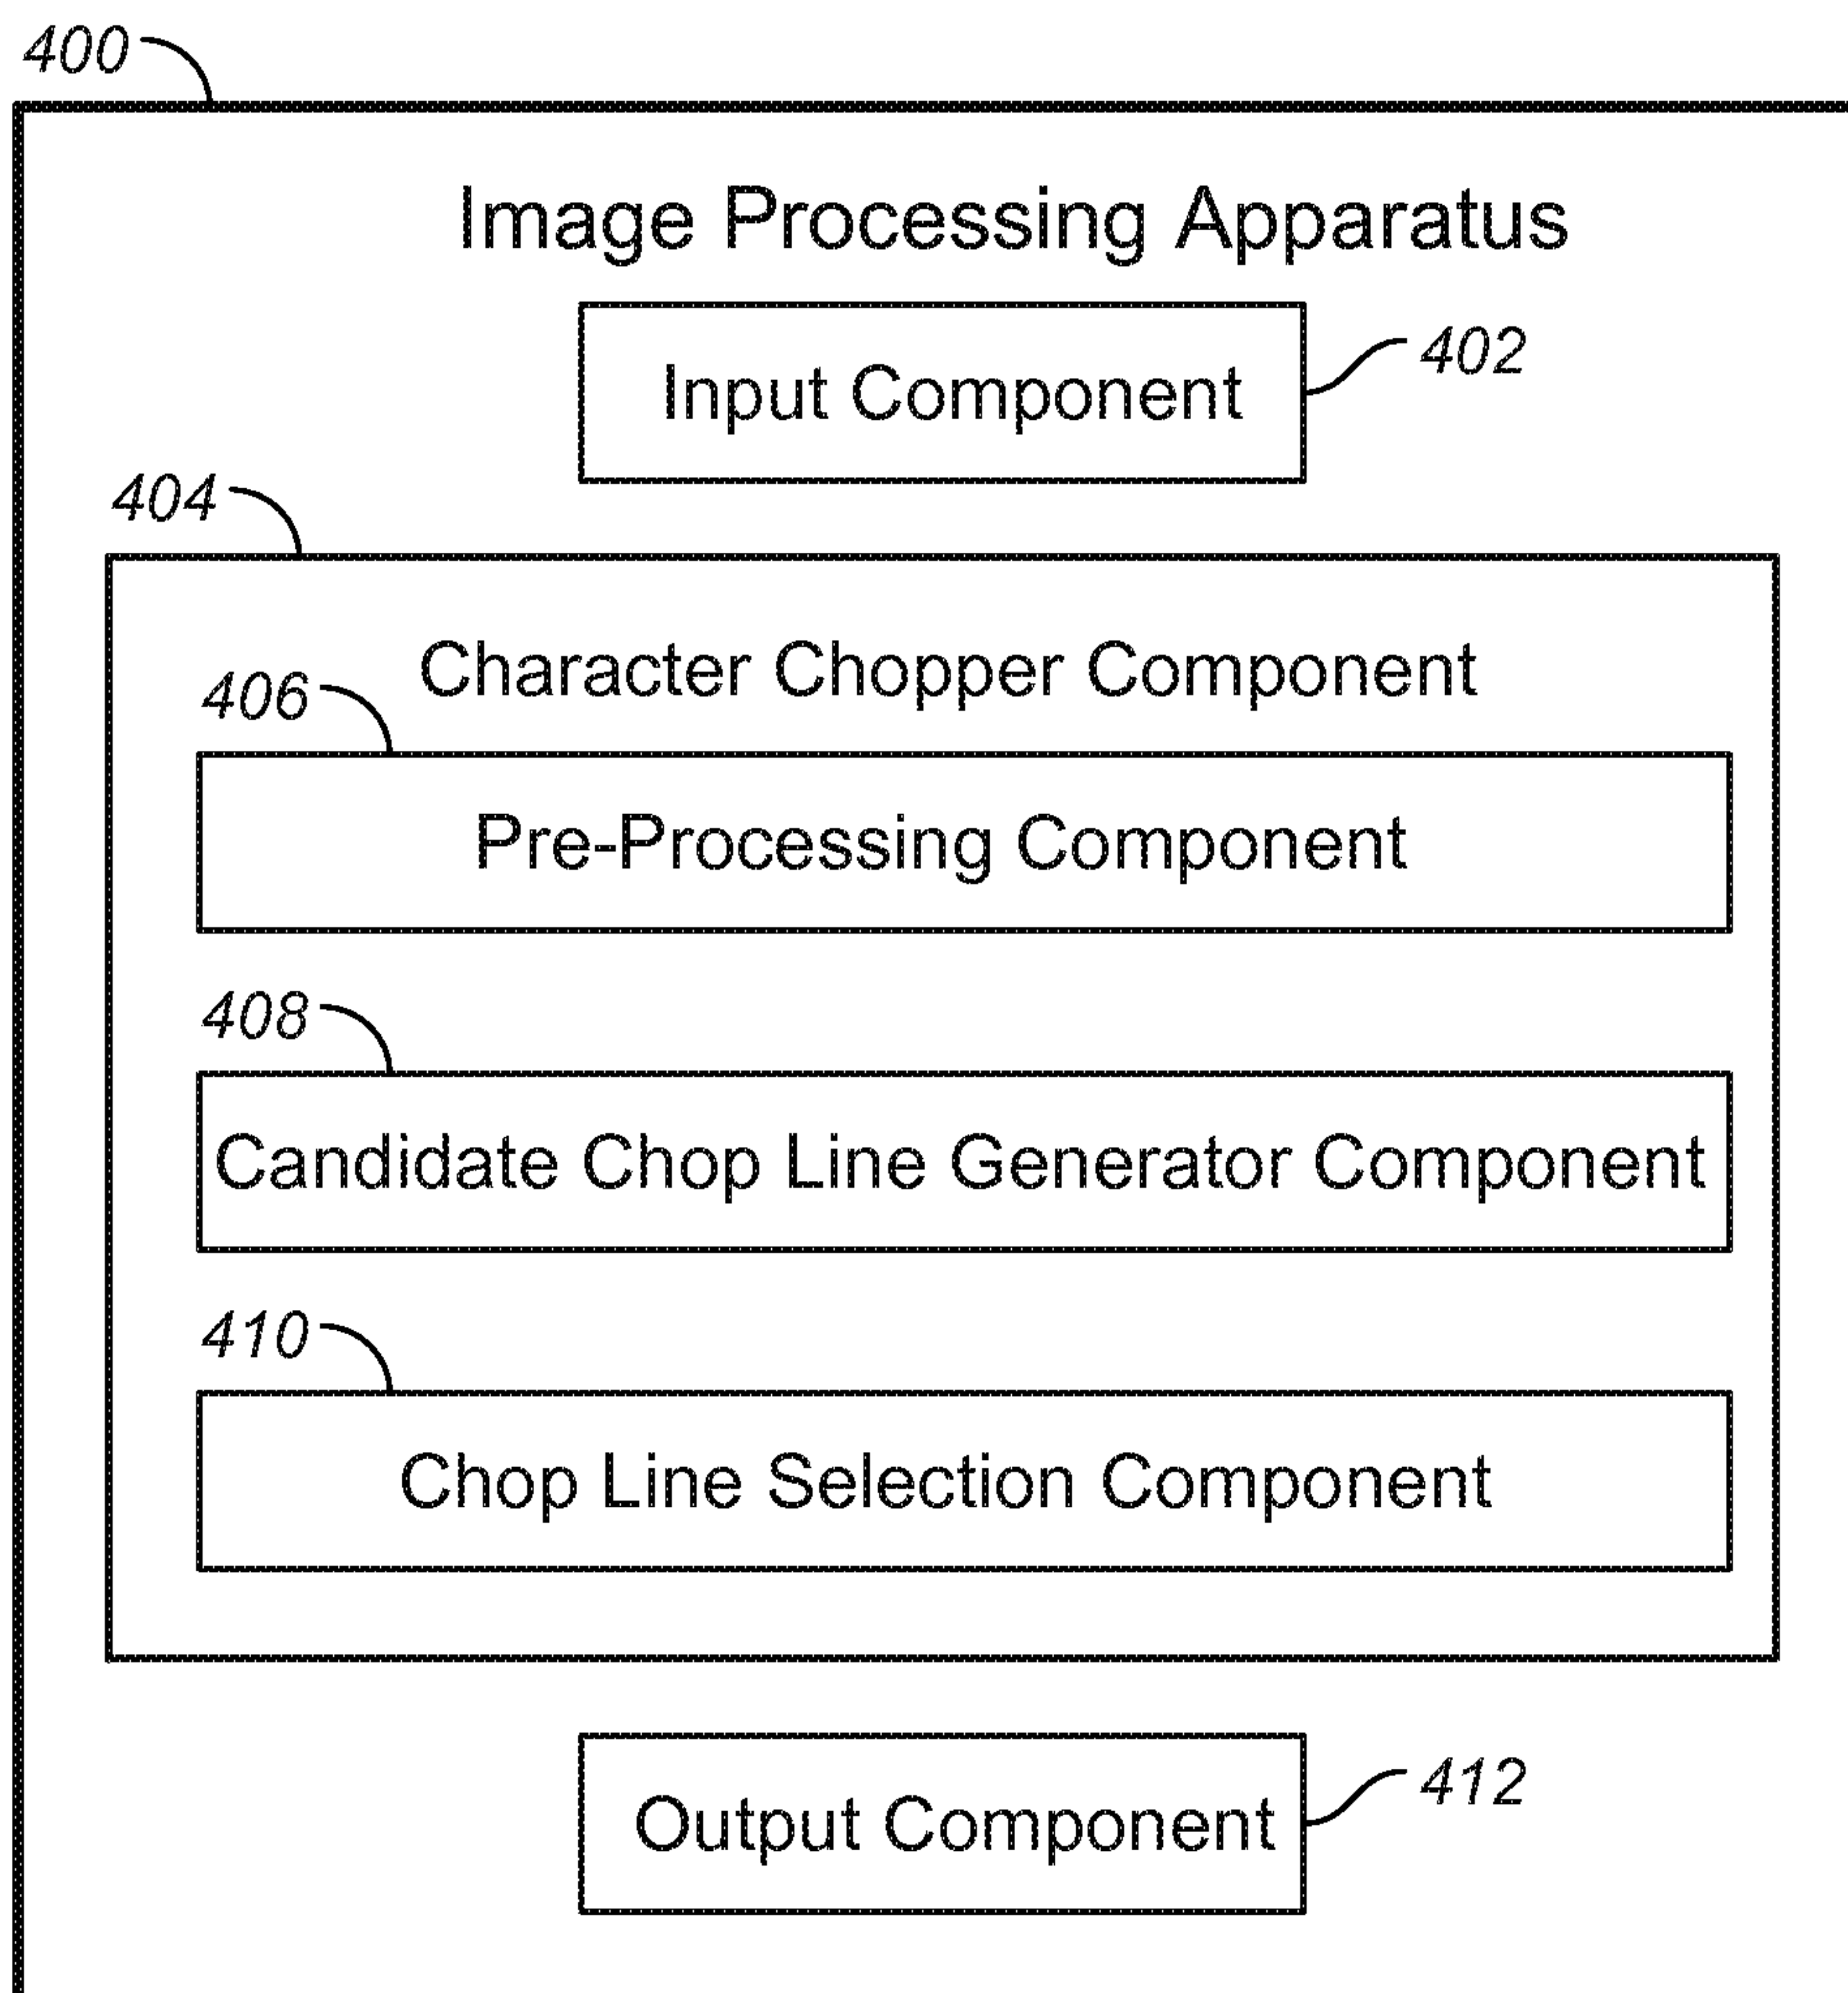


FIG. 4

3/4




FIG. 5



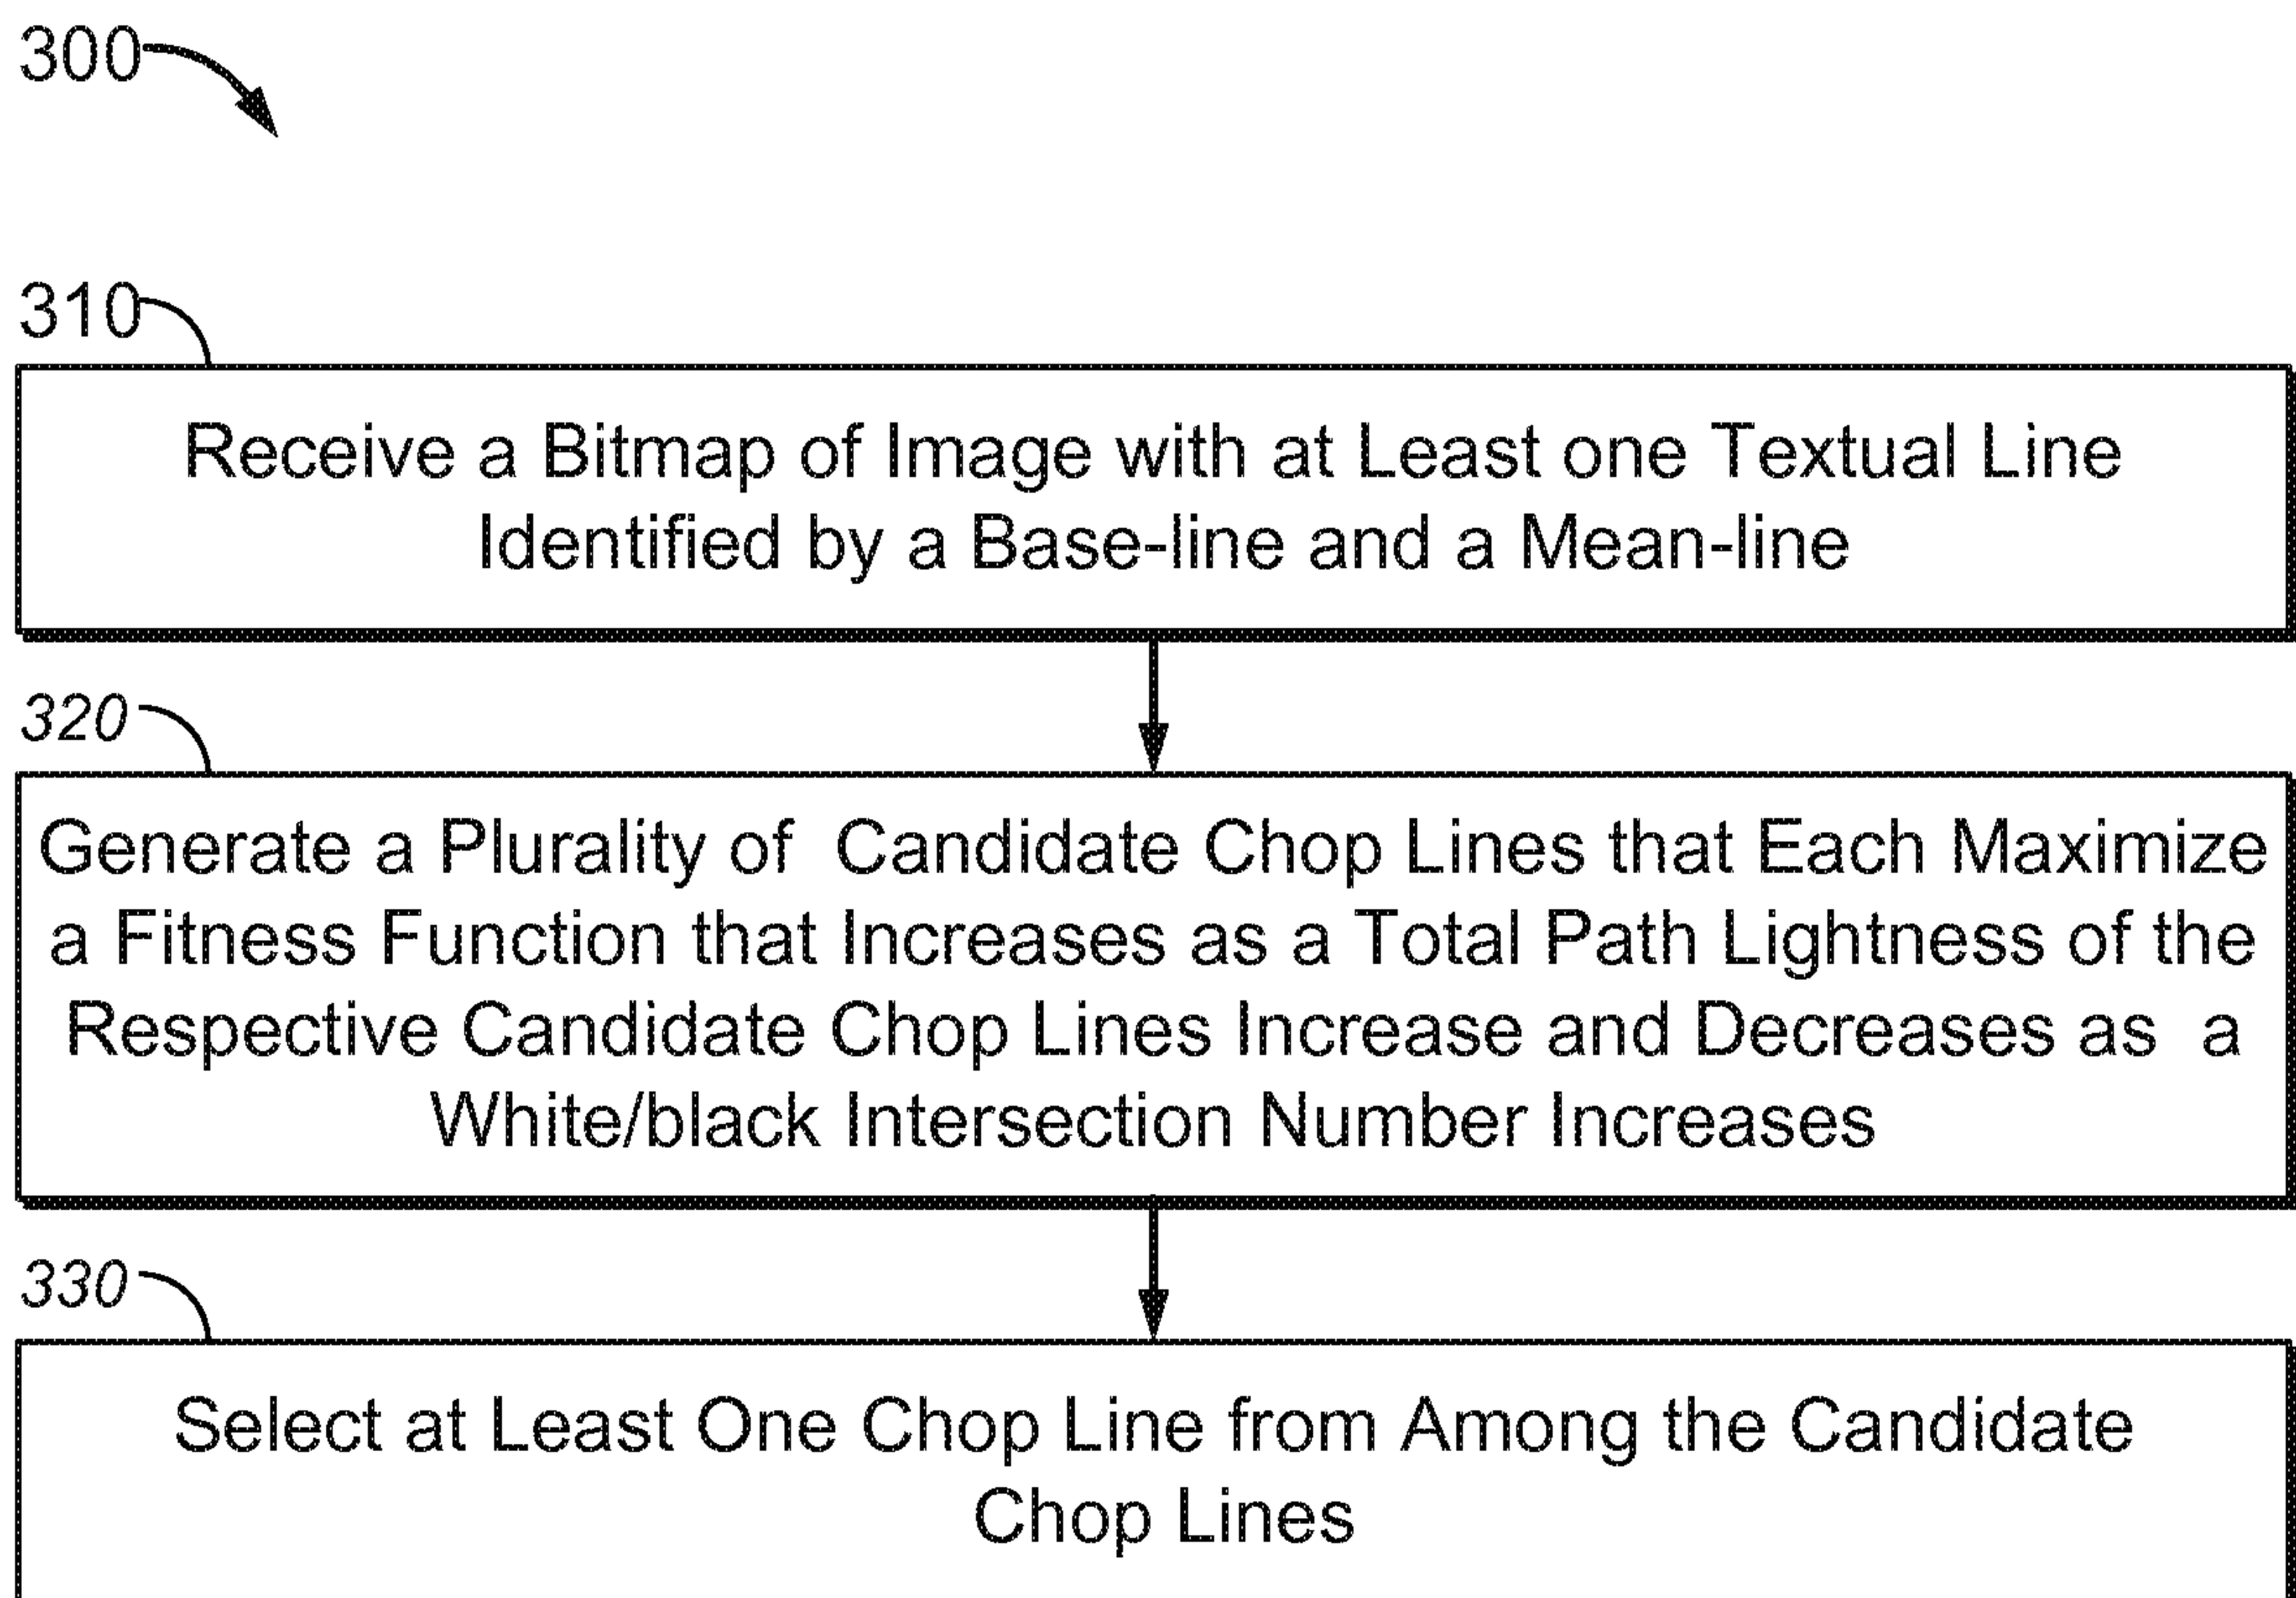


FIG. 6

4/4

**FIG. 7****FIG. 8**

400

Image Processing Apparatus

Input Component

402

404

Character Chopper Component

406

Pre-Processing Component

408

Candidate Chop Line Generator Component

410

Chop Line Selection Component

Output Component

412