




(86) Date de dépôt PCT/PCT Filing Date: 2011/03/07
(87) Date publication PCT/PCT Publication Date: 2011/09/15
(45) Date de délivrance/Issue Date: 2017/06/20
(85) Entrée phase nationale/National Entry: 2012/08/16
(86) N° demande PCT/PCT Application No.: US 2011/027365
(87) N° publication PCT/PCT Publication No.: 2011/112497
(30) Priorité/Priority: 2010/03/09 (US12/719,894)

(51) Cl.Int./Int.Cl. *G06K 9/20* (2006.01),
G06K 19/00 (2006.01)
(72) Inventeurs/Inventors:
NIJEMCEVIC, DJORDJE, US;
VUGDELIJA, MILAN, US;
DRESEVIC, BODIN, US
(73) Propriétaire/Owner:
MICROSOFT TECHNOLOGY LICENSING, LLC, US
(74) Agent: SMART & BIGGAR

(54) Titre : AJUSTEMENT DE LA RESOLUTION D'UNE IMAGE QUI COMPREND DU TEXTE SOUMIS A UN PROCEDE ROC

(54) Title: RESOLUTION ADJUSTMENT OF AN IMAGE THAT INCLUDES TEXT UNDERGOING AN OCR PROCESS



(57) **Abrégé/Abstract:**

An optical character recognition process characterizes text lines in a textual image by their base-line, mean-line and x-height. The base-line for at least one text line in the image is determined by finding a parametric curve that maximizes a first fitness function that depends on the values of pixels through which the parametric curve passes and pixels below the parametric curve. The mean-line is determined by incrementally shifting the base-line upward by predetermined amounts (e.g., a single pixel) until a second fitness function for the shifted base-line is maximized. The second fitness function is essentially the inverse of the first fitness function. Specifically, the second fitness function increases with increasing lightness of pixels immediately above the shifted base-line while also increasing with decreasing lightness of pixels through which the shifted base-line passes.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
15 September 2011 (15.09.2011)



(10) International Publication Number
WO 2011/112497 A3

(51) International Patent Classification:
G06K 9/20 (2006.01) *G06K 19/00* (2006.01)

(21) International Application Number:
PCT/US2011/027365

(22) International Filing Date:
7 March 2011 (07.03.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/719,894 9 March 2010 (09.03.2010) US

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

(72) Inventors: **NIJEMCEVIC, Djordje**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **VUGDELIJA, Milan**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **DRESEVIC, Bodin**; 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).

Published:

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

(88) Date of publication of the international search report:
17 November 2011

(54) Title: RESOLUTION ADJUSTMENT OF AN IMAGE THAT INCLUDES TEXT UNDERGOING AN OCR PROCESS




FIG. 5

(57) Abstract: An optical character recognition process characterizes text lines in a textual image by their base-line, mean-line and x-height. The base-line for at least one text line in the image is determined by finding a parametric curve that maximizes a first fitness function that depends on the values of pixels through which the parametric curve passes and pixels below the parametric curve. The mean-line is determined by incrementally shifting the base-line upward by predetermined amounts (e.g., a single pixel) until a second fitness function for the shifted base-line is maximized. The second fitness function is essentially the inverse of the first fitness function. Specifically, the second fitness function increases with increasing lightness of pixels immediately above the shifted base-line while also increasing with decreasing lightness of pixels through which the shifted base-line passes.

RESOLUTION ADJUSTMENT OF AN IMAGE THAT INCLUDES TEXT UNDERGOING AN OCR PROCESS

Background

[0001] Optical character recognition (OCR) is a computer-based translation of an image of
5 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
10 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
15 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 scanning
process is good, the performance of the OCR engine may be degraded when a relatively
20 low quality page of text is being scanned.

Summary

[0002] Optical character recognition requires the identification of the text lines in the
textual image in order to identify individual words and characters. The text lines can be
characterized by their base-line, mean-line and x-height. Determining these features may
25 be become difficult when the text lines are not perfectly horizontal, which may arise when
scanning some classes of documents (for example a thick book) in which the image suffers
from non-linear distortions. In such a case, the base-line and mean-line may not be
constant over an entire text line.

[0003] To overcome these problems, in one implementation the base-line for at least one
30 text line in the image is determined by finding a parametric curve that maximizes a first
fitness function that depends on the values of pixels through which the parametric curve
passes and pixels below the parametric curve. The base-line corresponds to the parametric
curve for which the first fitness function is maximized. The first fitness function is
designed so that it increases with increasing lightness or brightness of pixels immediately

51331-1251

below the parametric curve while also increasing with decreasing lightness of pixels through which the parametric curve passes.

[0004] In some implementations the mean-line can be determined by incrementally shifting the base-line upward by predetermined amounts (e.g., a single pixel) until a second
 5 fitness function for the shifted base-line is maximized. The second fitness function is essentially the inverse of the first fitness function. Specifically, the second fitness function increases with increasing lightness of pixels immediately above the shifted base-line while also increasing with decreasing lightness of pixels through which the shifted base-line passes.

[0005] In some implementations the x-height can be determined from the base-line
 10 and the mean-line which have already been calculated. In particular, the x-height is equal to the predetermined amount by which the base-line is shifted upward in order to maximize the second fitness function.

[0005a] According to one aspect of the present invention, there is provided a system that extracts information which characterized text lines in an image, comprising: an input
 15 component for receiving a bitmap of an input image that includes text lines; and a parameterizing engine that determines a base-line for at least one text line in the image by finding a parametric curve that maximizes a fitness function that depends on values of pixels through which the parametric curve passes and pixels below the parametric curve, wherein the base-line corresponds to the parametric curve for which the fitness function is maximized,
 20 wherein the fitness function is defined as $\text{fitness}(\text{baseline})$ and is equal to

$$\text{fitness}(\text{baseline}) = \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x] + 1, x] - \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x], x]$$

where: x and y are horizontal and vertical pixel coordinates respectively; $\text{img}[y, x]$ is a pixel value of the bitmap at location (y, x); width is a width of bitmap input image; and $\text{baseline}[x]$ is a y-coordinate of the base-line at position x.

51331-1251

[0005b] According to another aspect of the present invention, there is provided a method for extracting information which characterized text lines in an image, comprising: receiving a bitmap of an input image that includes text lines; and determining a base-line for at least one text line in the image by finding a parametric curve that maximizes a fitness function that depends on values of pixels through which the parametric curve passes and pixels below the parametric curve, wherein the base-line corresponds to the parametric curve for which the fitness function is maximized, wherein the fitness function increases with increasing lightness of pixels immediately below the parametric curve and also increases with decreasing lightness of pixels through which the parametric curve passes and is defined as fitness(baseline) and is equal to

$$fitness(baseline) = \sum_{x=0}^{width-1} img[baseline[x] + 1, x] - \sum_{x=0}^{width-1} img[baseline[x], x]$$

where x and y are horizontal and vertical pixel coordinates respectively; img[y, x] is a pixel value of the bitmap at location (y, x); width is a width of bitmap input image; and baseline[x] is a y-coordinate of the base-line at position x.

[0005c] According to still 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 as described above or detailed below.

[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.

Brief Description of the Drawings

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

51331-1251

[0008] FIG. 2 shows an example of a text-line in a scanned image which is not perfectly horizontal.

[0009] FIG. 3 illustrates the base-line for the text-line of a single word in a scanned image.

5 [0010] FIG. 4 is a flowchart illustrating a process of determining the x-height for different groups of text lines.

[0011] FIG. 5 shows one example of an image processing apparatus that may perform the process of extracting information concerning the text-lines in a textual image.

Detailed Description

10 [0012] FIG. 1 shows one illustrative example of a system 5 for optical character recognition (OCR) in an image which 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
15 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.

5 [0013] The OCR engine 20 receives a textual image as a bitmap of text lines. Three parameters of those text lines that need to be determined are the “base-line,” “mean-line,” and the “x-height.” 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). Second, the “mean-line” is defined as a horizontal line which passes through the top ends
10 of a majority of the characters in a line of text (excluding ascenders). Third, 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).

[0014] Knowing the precise base-line and x-height is important for a number of reasons,
15 particularly in differentiating between capital and lowercase letters of the same shape. If a text-line is perfectly horizontal and contains only one font style and size, base-line and x-height will hold a constant value over the entire line. Computing these values for a perfectly horizontal text-line is not a difficult task. However, when scanning some classes of documents (for example a thick book), the document image can suffer from non-linear
20 distortions. In such a case, the base-line coordinate is not going to be constant over an entire text line.

[0015] An example of a text line containing this artifact is shown in the FIG. 2. It can be seen that the text has a “wavy” appearance, which is caused by the decrease in the average letter position from the line’s middle towards the left or right. Artifacts of this nature make
25 it more difficult to determine the base-line.

[0016] Extracting x-height information from a textual image can also be problematic. For instance, sometimes a majority of a text line (or even an entire text line) is composed of capital letters or numbers. In such a case, extracting the x-height using the line’s bitmap as a unique information source is not reliable. FIG. 2 also shows the base-line, mean-line and
30 x-height.

[0017] As detailed below, a method is provided to compute the base-line of a deformed text-line in the form of a parametric curve. Moreover, the most probable x-height value of a given line is estimated using context information obtained from the entire image.

Base-Line Computation

[0018] At the outset, two observations can be made from the base-line's definition:

- Due to the nature of most fonts, the base-line will overlap with a significant amount of dark pixels originating from letter bottoms.
- 5 • Immediately below the base-line there are no dark pixels (except for descending letter parts).

Regardless of whether or not the base-line is strictly horizontal or (in case of non-linear deformations) "wavy," it should be possible to establish a simple fitness function based on at least two properties obtained from these observations.

- 10 • Property 1: As the pixels immediately below the baseline become lighter (i.e., brighter), the value of the fitness function will increase (and vice versa).
- Property 2: As the pixels overlapping with the baseline become darker, the value of the fitness function will increase (and vice versa).

The goal of finding the base-line in a given text-line bitmap translates to the
15 problem of finding a (curved) line with a maximal fitness function value.

A rasterized base-line can be implemented as an array: for each x-coordinate of an input bitmap, there should be one and only y-coordinate describing the local baseline value. Keeping this in mind, a simple proposal for the fitness function is:

$$fitness(baseline) = \sum_{x=0}^{width-1} img[baseline[x] + 1, x] - \sum_{x=0}^{width-1} img[baseline[x], x]$$

Where:

- 20 • x and y are horizontal and vertical pixel coordinates respectively (with the origin is at top-left corner)
- $img[y, x]$ is the input bitmap's pixel value at location (y, x)
- $width$ is the input bitmap's width
- $baseline[x]$ is a y-coordinate of the base-line at position x

25 [0019] It can be observed that the formula for the fitness function satisfies both Property 1 and Property 2. Since the pixel colors in a typical gray-scale image vary from black (value: 0) to white (value: 255), the following will hold true:

- As the pixels immediately below the baseline become lighter, the first addend in the formula will become larger.
- 30 • As the pixels on the baseline become darker, the second addend in the formula will become larger.

[0020] A simple diagram of text illustrating this idea is presented in FIG. 3. It can be observed from FIG. 3 that the base-line overlaps with a relatively large number of dark pixels, while pixels immediately below the base-line are completely white.

5 [0021] After defining the criteria that the baseline should fulfill, another question that arises is how “fast” the baseline should change across the text-line width when maximizing the fitness function. Clearly, this rate of change should be sufficient to track the line’s “waviness”.

[0022] On the other hand, the rate of change should not be too fast, because it is not desirable for the bottoms of descending characters to affect the base-line shape. One way
10 to address this issue is to define a base-line candidate through a small set of control parameters, and limiting the range of values each parameter can take. In this way the shape of the base-line candidate can be changed by changing its control parameters.

[0023] A curve maximizing the fitness function can be parameterized by defining it through a set of control points connected with straight line segments. The curve’s shape
15 can be varied by moving its control points. One way to control the movement of the control point in a manner that achieves good performance results is to only allow the control point to have freedom of movement in the vertical direction. This approach has shown that a set of 4-6 equidistant control points does a good job in modeling a common “wavy” baseline.

20 [0024] A second way of parameterizing the curve for the base-line is by defining it as a B-spline. Changing its shape can be done by varying the spline coefficients.

[0025] In general, finding the exact shape that maximizes some fitness function can be thought of as a classical optimization problem which can be solved using well-known techniques. Depending on the nature and number of parameters used to describe the base-
25 line curve, a genetic search, dynamic programming, or some other technique can be used.

[0026] If a genetic search is performed, an initial population can be a set of curves with parameters randomly set within some reasonable range. New offspring can be formed by taking two high-fitness curves and mixing their parameters into a new curve. Mutation can be done by slightly varying curve parameters.

30 [0027] The curve parameters can be optimized by dynamic programming as well. The solution requires finding an optimal path starting at the text-line’s left side and moving towards its right side, while obeying the spatial constraints imposed by the common curve shape.

X-height computation

[0028] The mean-line (a line determining where non-ascending lowercase letters terminate) can be computed in a way quite similar to the base-line computation procedure described above. Actually, it is enough to invert the fitness function described above and re-run the algorithm. That is, the fitness function for the mean-line should satisfy the following two properties:

- Property 1: As the pixels immediately above the baseline become lighter, the value of the fitness function will increase (and vice versa).
- Property 2: As the pixels overlapping with the baseline become darker, the value of the fitness function will increase (and vice versa).

[0029] Once the mean-line is determined, the x-height can then be extracted by simply subtracting the corresponding mean-line and base-line coordinates. However, this process introduces an additional computational load, effectively doubling the entire feature extraction execution time.

[0030] In practice, non-linear deformations of the type discussed herein have no influence on individual letter dimensions. In other words, the x-height does not change across the “wavy” text-line, provided that the line contains letters of the same font style and size. This conclusion facilitates the process of computing the x-height since it directly implies that the curves for the mean-line and the base-line will have exactly the same shape.

Accordingly, the mean-line can be computed in the following way: the curve for the base-line is shifted pixel by pixel towards the text-line’s top, and the inverted fitness function is computed each time the curve is shifted upward. The shifted curve that results when the fitness function reaches its maximum value will be the mean-line. The number of pixels by which the base-line curve is shifted upward to obtain the mean-line is equal to the x-height.

[0031] Sometimes an input bitmap of an individual text-line cannot be used as the only source of information to obtain a single value for x-height over an entire image. For example, some text lines may be short lines composed of numbers only. Another example is a caption that is in all capital letters. Because of such cases, the x-height computation may sometimes be performed in a somewhat more sophisticated manner.

[0032] In this implementation, before computing the x-height, it is determined whether the text-lines in the images should be divided into different groups that are each likely to contain text-lines with different x-heights. Such text-line groups may be determined in a

51331-1251

variety of different ways. For instance, text-lines may be grouped according to their dominant letter stroke width. This approach essentially assumes that different x-heights arise from the use of different fonts and font sizes and that each such font and font size is characterized by a different dominant stroke width. Thus, groups of text lines with a

5 common dominant stroke width likely have a common x-height.

[0033] The dominant stroke width may be determined at this stage of the OCR process or it may have been determined in an earlier stage of processing which precedes the text line analysis described herein. One example of a method for determining stroke width is shown in U.S. Patent Appl. Serial No. 12/720,732.

[0034] In one alternative, instead of grouping text lines by their dominant stroke width, individual words can be divided into their own groups.

[0035] To determine the x-height of a particular group, begin by defining a mean-line candidate [j] as a base-line's curve shifted by j pixels up. Next, for each group, a common
15 buffer is established. For each text-line in a group, an inverted fitness function of the mean-line candidate [j] is added to the buffer. At the end of the process, the buffer's element j will contain a sum of inverted fitness functions for all the mean-line candidates [j] within the particular group. The most-probable x-height value for the particular group corresponds to the value of j for which the buffer has its maximum value.

20 [0036] A flowchart illustrating the process of determining the x-height for different groups of text lines is shown in FIG. 4. The process begins in step 105 when the text-lines in an image are divided into groups by any appropriate criterion such as font size, dominant stroke width, or the like. For each group, the process continues from step 110 to step 115 in which an accumulation buffer is established and initialized to a value of zero. Next, for
25 each text line within the group the process proceeds from step 120 to step 125 in which j is initialized to zero and the mean-line candidate is initialized to the base-line. The value of j is incremented by 1 in step 135, which corresponds to shifting the base-line curve upwards by one pixel. The fitness function for this mean-line (which corresponds to the inverse fitness function of the base-line) is calculated in step 140. Also in step 140, the
30 accumulation buffer is defined as the sum of its previous value and the value of the fitness function that has just been calculated. Decision step 145 then determines if the maximum value of the accumulation buffer has been reached. If so, then in step 150 the current value of j corresponding to this maximum value is determined to be the x-height for this group. Alternatively, if the maximum value of the accumulation buffer has not been reached, the

process proceeds from decision step 145 back to step 130 in which the current mean-line is shifted upward by 1 pixel. This process continues until the maximum value of the accumulation buffer has been reached. Once the x-height value has been determined for this group the process returns to step 120 and repeats for any remaining groups of text

5 lines, finally ending at step 155.

[0037] FIG. 5 shows one example of an image processing apparatus 300 that may perform the process of extracting information concerning the text-lines in a textual image. The apparatus, which may be incorporated in an OCR engine, can be used by the OCR engine to determine the base-line, mean-line and x-height of the text lines in the image. The

10 apparatus includes an input component 302 for receiving an input image and a parameterizing engine 310 for finding parametric curves that correspond to the base-line and the mean-line of the text lines in the image. The parameterizing engine 310 includes a base-line determination component 322, a mean-line determination component 324 and an x-height determination component 324. The apparatus 300 also includes an output
15 component 330 that generates the information concerning the text-lines in a form that allows it to be employed by subsequent components of the OCR engine.

[0038] 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
20 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
25 computer and/or distributed between two or more computers.

[0039] 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
30 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,

51331-1251

those skilled in the art will recognize many modifications may be made to this configuration without departing from the scope of the claimed subject matter.

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 the

5 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.

51331-1251

CLAIMS:

1. A system that extracts information which characterized text lines in an image, comprising:

an input component for receiving a bitmap of an input image that includes text
5 lines; and

a parameterizing engine that determines a base-line for at least one text line in the image by finding a parametric curve that maximizes a fitness function that depends on values of pixels through which the parametric curve passes and pixels below the parametric curve, wherein the base-line corresponds to the parametric curve for which the fitness
10 function is maximized, wherein the fitness function is defined as $\text{fitness}(\text{baseline})$ and is equal to

$$\text{fitness}(\text{baseline}) = \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x] + 1, x] - \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x], x]$$

where:

x and y are horizontal and vertical pixel coordinates respectively;

$\text{img}[y, x]$ is a pixel value of the bitmap at location (y, x);

15 width is a width of bitmap input image; and

$\text{baseline}[x]$ is a y-coordinate of the base-line at position x.

2. The system of claim 1 wherein at least one control parameter constrains at least one feature of the parametric curve.

3. The system of claim 2 wherein the feature of the parametric curve that is
20 determined by the control parameter is a maximum rate of change of the parametric curve along the text line.

51331-1251

4. The system of claim 1 wherein the parametric curve includes a plurality of control points connected by straight lines, wherein the control points are constrained to move only in a vertical direction.

5. The system of claim 1 wherein the parametric curve is defined as a B-spline
5 having a shape that is determined by its spline coefficients.

6. The system of claim 1 further comprising maximizing the fitness function using an optimization technique is selected from the group consisting of a genetic search and dynamic programming.

7. The system of claim 1 wherein the parameterizing engine further comprises a
10 mean-line determination component that determines a mean-line for the at least one line of text.

8. The system of claim 7 wherein the mean-line determination component determines the mean-line by maximizing a second fitness function for a second parametric curve, wherein the second fitness function increases with increasing lightness of pixels
15 immediately above the second parametric curve and also increases with decreasing lightness of pixels through which the second parametric curve passes.

9. The system of claim 8 wherein the mean-line determination component determines the mean-line by:

incrementally shifting the base-line upward by predetermined amounts until a
20 second fitness function for the shifted base-line is maximized, wherein the second fitness function increases with increasing lightness of pixels immediately above the shifted base-line and also increases with decreasing lightness of pixels through which the shifted base-line passes.

10. The system of claim 9 further comprising an x-height determination
25 component that determines an x-height for the at least one text line, wherein the x-height is

51331-1251

equal to a predetermined amount by which the base-line is shifted upward in order to maximize the second fitness function.

11. The system of claim 1 wherein the parameterizing engine determines a different x-height for different groups of text-lines in the input image.

5 12. The system of claim 11 wherein the parameterizing engine divides the text-lines in the input image into groups based on their dominant stroke-width.

13. The system of claim 1 wherein the parameterizing engine determines a most probable x-height for different groups of text-lines in the input image based on an inverted fitness function of a mean-line candidate for each text-line in the group.

10 14. A method for extracting information which characterized text lines in an image, comprising:

receiving a bitmap of an input image that includes text lines; and

determining a base-line for at least one text line in the image by finding a parametric curve that maximizes a fitness function that depends on values of pixels through
 15 which the parametric curve passes and pixels below the parametric curve, wherein the base-line corresponds to the parametric curve for which the fitness function is maximized, wherein the fitness function increases with increasing lightness of pixels immediately below the parametric curve and also increases with decreasing lightness of pixels through which the parametric curve passes and is defined as $\text{fitness}(\text{baseline})$ and is equal to

$$\text{fitness}(\text{baseline}) = \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x] + 1, x] - \sum_{x=0}^{\text{width}-1} \text{img}[\text{baseline}[x], x]$$

20 where

x and y are horizontal and vertical pixel coordinates respectively;

$\text{img}[y, x]$ is a pixel value of the bitmap at location (y, x);

51331-1251

width is a width of bitmap input image; and

baseline[x] is a y-coordinate of the base-line at position x.

15. The method of claim 14 further comprising maximizing the fitness function using an optimization technique selected from the group consisting of a genetic search and
5 dynamic programming.

16. The method of claim 14 further comprising determining a mean-line for the at least one line of text by incrementally shifting the base-line upward by predetermined amounts until a second fitness function for the shifted base-line is maximized, wherein the second fitness function increases with increasing lightness of pixels immediately above the
10 shifted base-line and also increases with decreasing lightness of pixels through which the shifted base-line passes.

17. The method of claim 16 further comprising determining an x-height for the at least one text line, wherein the x-height is equal to a sum of the predetermined amounts by which the base-lines is shifted upward in order to maximize the second fitness function.

15 18. A non-transitory computer-readable medium, having stored thereon, computer executable instructions, that when executed, perform a method according to any one of claims 14 to 17.

1/3




FIG. 1

1871, p. 278 f.). Cp. Bunyiu Nanjio, *A Catalogue of the Chinese Translation*

FIG. 2




FIG. 3

2/3




FIG. 4

3/3




FIG. 5

Image Processing Apparatus

310

Input Component

320

Parameterizing Engine

322

Base-line Determination Component

324

In Each-line Determination Component

326

X-height Determination Component

330

Output Component