



(22) Date de dépôt/Filing Date: 2002/09/03

(41) Mise à la disp. pub./Open to Public Insp.: 2003/03/10

(30) Priorité/Priority: 2001/09/10 (09/950,205) US

(51) Cl.Int.⁷/Int.Cl.⁷ G06K 9/00, G06F 17/00


(71) Demandeur/Applicant:
XEROX CORPORATION, US

(72) Inventeurs/Inventors:
BERN, MARSHALL W., US;
GOLDBERG, DAVID, US;
YE, MING, US

(74) Agent: SIM & MCBURNEY

(54) Titre : METHODE DE RAPPROCHEMENT DE DOCUMENTS ET DE DUPLICATION D'ANNOTATIONS

(54) Title: DOCUMENT MATCHING AND ANNOTATION LIFTING



(57) Abrégé/Abstract:

A method for matching an original document image with a copy image is disclosed. The original document image is defined as an ensemble of blocks, each of the blocks having neighboring blocks. A mask is formed for each of the blocks by dilating foreground pixels included in the block. A best match is searched in the copy image, for each of the blocks, using a search window, by correlating pixels in the mask with pixels in a corresponding portion in the copy image. Each of the best matches has a matching score. Each of the blocks is indicated as "unmatched" when the corresponding matching score is less than a matching threshold. A displacement vector is computed for each of the matched blocks.



ABSTRACT

A method for matching an original document image with a copy image is disclosed. The original document image is defined as an ensemble of blocks, each of the blocks having neighboring blocks. A mask is formed for each of the blocks by dilating foreground pixels included in the block. A best match is searched in the copy image, for each of the blocks, using a search window, by correlating pixels in the mask with pixels in a corresponding portion in the copy image. Each of the best matches has a matching score. Each of the blocks is indicated as "unmatched" when the corresponding matching score is less than a matching threshold. A displacement vector is computed for each of the matched blocks.

DOCUMENT MATCHING AND ANNOTATION LIFTING

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

The present invention relates generally to methods and systems for image
5 processing, and particularly to methods and systems for matching documents and
lifting annotations present in only one of the two matched documents.

DESCRIPTION OF RELATED ART

Traditionally many information gathering tasks have been carried out using
paper documents. Examples include opinion surveys, tax returns, check processing,
10 and revising and editing text. An important step in many such tasks is annotation
lifting, that is, the problem of extracting content added to an original document. A
special case of annotation lifting is called form dropout, in which the added content is
assumed to appear only in certain predefined locations. In order to extract
annotations from an annotated document, the annotated document has to be
15 matched with an image of the original document of which the annotated document is
a copy.

Matching two images of a same document is not a simple task since the two
images could differ significantly due to faxing, scanner distortions, or degradation
through multigeneration copying. Preprinted data of an original document image
20 may appear distorted when compared to preprinted data of an annotated copy of the
same original document, and the distortion may be different for different parts of the
image.

There is a need for an efficient method for registering two such images and
for separating out annotations. There is also a need for an efficient algorithm for
25 detecting and repairing broken strokes in the annotations, which occur when the
annotations cross or touch preprinted data.

In addition to annotation lifting, such method for document matching also finds
application in duplicate removal and in document image authentication in which the
task is to confirm that a document has not changed since authorship.

SUMMARY OF THE INVENTION

A method for matching an original document image with a copy image is disclosed. The original document image is defined as an ensemble of blocks, each of the blocks having neighboring blocks. A mask is formed for each of the blocks by dilating foreground pixels included in the block. A best match is searched in the copy image, for each of the blocks, using a search window, by correlating pixels in the mask with pixels in a corresponding portion in the copy image. Each of the best matches has a matching score. Each of the blocks is indicated as "unmatched" when the corresponding matching score is less than a matching threshold. A displacement vector is computed for each of the matched blocks.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the present invention will become apparent from the following detailed description of the present invention in which:

FIG. 1 is a block diagram of one embodiment of the system of the present invention.

FIG. 2 is a flowchart illustrating the process of the global alignment module 110 (FIG. 1) as implemented in one embodiment of the present invention.

FIG. 3 is a flowchart illustrating the process of the document matching module 120 (FIG. 1) as implemented in one embodiment of the present invention.

FIG. 4 is a flowchart illustrating an exemplary hierarchical search process as implemented in one embodiment.

FIG. 5 illustrates an example of a mask obtained by dilation of the black pixels in a block of the original document.

FIG. 6 illustrates a hierarchical search covering a 15 x 15 search window.

FIG. 7A shows the initial distortion map of a test document after the initial matching.

FIG. 7B shows the updated distortion map.

FIG. 7C shows the final distortion map.

FIG. 8 is a flowchart illustrating an exemplary registration process of the registration module 130 (FIG. 1).

FIG. 9 is a flowchart illustrating an exemplary authentication process of the authentication module 150 (FIG. 1).

FIG. 10 is a flowchart of an exemplary process 1000 of the initial annotation image module 142 (FIG. 1).

5 FIG. 11 is a flowchart of the process 1100 of the broken stroke repair module 144 (FIG. 1) as implemented in one embodiment of the present invention.

FIG. 12 illustrates how the broken stroke ends are defined.

FIG. 13A illustrates the case where a broken end may be chosen by more than one other broken end.

10 FIG. 13B illustrates the exceptional case where a broken end is itself a connected component in the initial annotation image *la*.

FIG. 14A illustrates the line tracing protocol.

FIG. 14B illustrates two broken end cases that are difficult to resolve.

FIG 15A shows the grayscale image of an annotated copy.

15 FIG. 15B shows the initial annotation image.

FIG. 15C shows the broken stroke ends.

FIG. 15D shows the final annotation image after broken stroke end reconstruction.

DETAILED DESCRIPTION OF THE INVENTION

The present invention discloses a method and a system for matching an
 20 original document image with a copy image. The original document image is defined as an ensemble of blocks, each of the blocks having neighboring blocks. A mask is formed for each of the blocks by dilating foreground pixels included in the block. A best match is searched in the copy image, for each of the blocks, using a search
 25 window, by correlating pixels in the mask with pixels in a corresponding portion in the copy image. Each of the best matches has a matching score. Each of the blocks is indicated as "unmatched" when the corresponding matching score is less than a matching threshold. A displacement vector is computed for each of the matched blocks.

The present invention also discloses a method for annotation lifting.

It will be obvious to one skilled in the art that the present invention can be used to match an original document image that is in either gray-level or binary form with a copy image that is in either gray-level or binary form.

5 Original document images sent over the Internet are likely in binary form to save memory and bandwidth. Scanned images stored by low-cost scanners are also usually in binary form. Therefore, the case where an original document image in binary form (to save memory and bandwidth) needs to be matched with a gray-level copy image will likely find more applications in the business environment. One
10 example of business applications is the validation of drug prescription forms.

The present invention will be described in detail as applied to the case where the original document image is in binary form and the copy image is in gray-level form. However, this is not to be construed as a limitation of the present invention.

The method of the present invention does not depend on assumptions of
15 document structure or annotation location. It can handle general document images containing text, drawings, and halftone images. Annotations can be arbitrarily placed, even touching or overlapping preprinted data (i.e., original content of the original document image).

The following terms are used throughout the description of the present
20 invention, and are defined here for clarity.

The term "connectivity" denotes the criterion that describes how pixels in an image form a connected group. For example, a connected component is "8-connected" if diagonally adjacent pixels are considered to be touching.

The term "foreground pixels" denotes pixels that form the text or image. In a
25 binary image, if the writing is black on a white background, then the foreground pixels are black. If the writing is white on a black background, then the foreground pixels are white.

The term "structuring element" denotes the matrix that is used to define a neighborhood shape and size for a morphological operation such as dilation.

30 The term "dilation" denotes the operation where the value of the output pixel is the maximum value of all the pixels in the input pixel's neighborhood. In a binary image, if any of the pixels is set to the value 1, the output pixel is set to 1.

In the following description of one embodiment of the present invention, the global alignment and registration processes are applied to the original document image. It is obvious that they can be applied to the copy image instead.

FIG. 1 is a block diagram of one embodiment of the system of the present invention. System 100 comprises a global alignment module 110, a document matching module 120, a registration module 130. Optionally, depending on an application of the present invention, system 100 can also include an annotation lifting module 140 or an authentication module 150.

The global alignment module 110 receives as inputs an original document image 101 and a gray-level copy 103, and outputs an aligned original document image 111 and a binary copy 113. The original document image is assumed to be in binary form, to save memory and bandwidth. If it is in a gray-level form, it can be easily thresholded to produce a binary scan.

The document matching module 120 matches the aligned original document image 111 with the binary copy 113 and produces a map of local displacement vectors 121. It also produces a final binary copy 123.

The registration module 130 applies the map of local displacement vectors 121 on the aligned original document image 111 to produce a registered original document image 131.

In an application where extraction of annotations present only on the gray-level copy is required, system 100 also includes the annotation lifting module 140. The annotation lifting module 140 receives as inputs the registered original document image 131 and the gray-level copy 103 and outputs an annotation image 145. The annotation lifting module 140 includes an initial annotation image module 142 and a broken stroke repair module 144. The initial annotation image module 142 produces an initial annotation image 143. The broken stroke repair module 144 repairs broken strokes, if any, in the initial annotation image 143 and outputs the annotation image 145.

In an application where authentication of the gray-level copy is required, system 100 also includes the authentication module 150. The authentication module 150 receives as inputs the registered original document image 131 and the final binary copy 123 and outputs an indicator 151 which indicates whether the gray-level copy is authenticated.

FIG. 2 is a flowchart illustrating the process of the global alignment module 110 (FIG. 1) as implemented in one embodiment of the present invention. Process 200 forms a first binary copy by binarizing the gray-level copy with a threshold (block 202). Process 200 computes a horizontal profile and a vertical profile for the original document image and for the first binary copy (block 204). The horizontal profile is computed by summing the columns of black pixels. The vertical profile is computed by summing the rows of black pixels. Then, process 200 computes the displacement in the x-coordinate by correlating the horizontal profile of the original document image with the horizontal profile of the first binary copy, and the displacement in the y-coordinate by correlating the vertical profile of the original document image with the vertical profile of the first binary copy (block 206). The displacements in the x-coordinate and y-coordinate, are computed from the maxima in the correlations of horizontal and vertical profiles. Finally, process 200 spatial transforms the original document image using the displacement in the x-coordinate and the displacement in the y-coordinate to produce an aligned original document image (block 208).

After the global alignment, the aligned original document image is divided into overlapping blocks for local matching. In one embodiment, each block has a size of 200 x 200 pixels with an overlap of 10 pixels with each neighboring block. The overlap prevents gaps between blocks in the aligned original. The result of the local matching is a distortion map in which each block is associated with a displacement vector and a matching score. The optimal distortion map is defined to be the one that maximizes block matching scores and minimizes variance of neighboring displacement vectors. This map is approximated through iterative relaxation, with initial matches given by a hierarchical search scheme. The hierarchical search scheme will be described later.

FIG. 3 is a flowchart illustrating the process of the document matching module 120 (FIG. 1) as implemented in one embodiment of the present invention.

Process 300 defines the aligned original document image as an ensemble of overlapping image blocks (block 302), then forms a mask for each of the image blocks by dilating the black pixels in the image block (block 304). For each image block, process 300 searches for a best match in the first binary copy, using a first search window, by correlating pixels in the mask with pixels in a corresponding portion in the first binary copy. Each of the best matches is associated with a

matching score. Each image block is marked as “unmatched” when the corresponding matching score is less than a matching threshold. Process 300 computes a displacement vector for each of the matched blocks (block 306). If all the image blocks are matched, process 300 terminates.

5 Otherwise, based on the matched blocks, process 300 computes a second binarization threshold for the gray-level copy (block 308). Using the second binarization threshold, process 300 binarizes the gray-level copy to produce a second binary copy (block 310). Process 300 assigns to each of the unmatched blocks a displacement vector computed from an average of respective displacement
10 vectors of the neighboring blocks of the unmatched block (block 312). In one embodiment, eight neighboring blocks are used in this computation. Process 300 searches for a new best match in the second binary copy, for each of the blocks, using a second search window, by correlating pixels in the mask with pixels in a corresponding portion in the second binary copy (block 314). The second search
15 window is can be much smaller than the first search window. Each of the new best matches is associated with a new matching score. Each of the image blocks is marked as “unmatched” when the corresponding new matching score is less than a matching threshold. A new displacement vector is computed for each of the new
20 matched blocks. If convergence is reached, that is, if all the image blocks are matched, process 300 terminates. Otherwise, process 300 repeats block 312 and block 314 until convergence is reached.

In block 306 and 314 of FIG. 3, process 300 uses a hierarchical search to find the best match of an image block.

FIG. 4 is a flowchart illustrating an exemplary hierarchical search process
25 400.

Process 400 defines a first box smaller than the search window and centered within the search window (block 402). Process 400 computes matching scores for matches located at points located within the first box. In one embodiment, it computes matching scores for 9 points which are the corners, side midpoints and
30 center of the first box (block 404). Process 400 designates the point having the highest matching score as the best match (block 406), defines a box smaller than the previous box and centered around the best match (block 408). Process 400 computes matching scores for matches located at points located at the corners, side

midpoints and center of the smaller box (block 410). Process 400 designates the point having the highest matching score as the best match (block 412). If the 9 points of the smaller box are not adjacent to each other, process 400 repeats blocks 408, 410 and 412. If the 9 points of the smaller box are adjacent to each other, then
 5 process 400 checks whether the best match is at the boundary of the search window (block 416). If it is not, process 400 terminates. Otherwise, process 400 defines a new search window centered around the best match (block 418) and repeats blocks 404 through 416.

To minimize the interference of possible annotations, the matching is
 10 examined only within a mask obtained by dilating the black pixels in the original. FIG. 5 illustrates an example of a mask obtained by dilation of the black pixels in a block of the original document image.

The matching score for a given displacement is defined as the number of overlapping black pixels divided by the average number of black pixels in the mask
 15 and the region of the binary copy corresponding to the mask. This score has range $[0,1]$, and allows comparison among different blocks on the same page.

If the block has more than 75% black pixels, it is assumed to be primarily white writing on a black background. The colors are then flipped before the mask is created and before the matching is started. Nearly pure black or white blocks are
 20 considered "empty" and their matches are initially disregarded. The displacement vectors for these "empty" blocks will be computed from neighboring matches.

In one embodiment, the distortion map is initialized by independently searching for each non-empty block's best match within a 29×29 window. Rather than using a brute-force implementation, which would compute all $841 = (29 \times 29)$
 25 matching scores, a hierarchical search scheme is used.

The search scheme first computes the matches for the corners, center, and side midpoints of a box, then matches for corners and side midpoints of a smaller box, centered at the highest-scoring point from the first level, and so forth. A 15×15 search window is covered by using box sizes of 9×9 , 5×5 , and 3×3 . If the best
 30 match turns out to be on the boundary, another such 15×15 search centered on the best match is performed, thus covering a 29×29 window overall. For the 29×29 window case, this hierarchical search scheme reduces the number of matching score computations to a worst case of 49. This represents a substantial reduction in

computations, compared to the 841 computations that would be required by the brute-force implementation. In general, if the brute-force scheme requires N computations, then this hierarchical scheme requires only $\log(N)$ computations.

FIG. 6 illustrates a hierarchical search covering a 15×15 search window.

5 Adaptive thresholding is used in the hierarchical search scheme. A match is accepted for a given block if its matching score exceeds a certain value. These good matches (even a very few will do) are then used to determine an optimal threshold for rebinarizing the gray-level copy (recall that the grayscale copy is originally binarized at a standard threshold.).

10 The optimal threshold is chosen to obtain the best possible agreement between the numbers of black pixels in the original document image and the rebinarized copy within the matched masks.

The whole grayscale copy, not just the matched blocks, is rebinarized using the optimal threshold to produce a final binary copy. The final binary copy replaces
15 the first binary copy for the rest of the search process.

After rebinarization, the distortion map is recomputed with the matching of the blocks to the final binary copy. The matching scores are updated and select blocks with high enough scores are selected as "good" ones.

FIG. 7A shows the initial distortion map of a test document image after the
20 initial matching. Brighter pixels indicate higher matching scores, and black ones indicate empty blocks.

The initial distortion map is improved through iterative relaxation. First, each not-yet well matched nonempty block is assigned a matching score that is computed using the average displacement of its good 8-connected neighbors (if any), and the
25 new assignment is kept if it brings a higher matching score. This process is iterated until convergence, and afterwards, 8-neighbor averages are propagated to empty blocks.

FIG. 7B shows the updated distortion map.

For each block, the variance of its 8 neighbors' displacement vectors is also
30 computed with respect to the block's own displacement vector. For each block with either a low matching score or a high variance, the average displacement of its good neighbors (if any) is taken as the starting position. The search process search for

the best match within a 7 x 7 window, and updates again upon obtaining a higher score. This process is iterated until convergence.

FIG. 7C shows the final distortion map.

Using the map of the displacement vectors produced by the document
5 matching module 120 (FIG. 1), the aligned original document image 111 (FIG.1) can now be registered.

FIG. 8 is a flowchart illustrating an exemplary registration process 800 of the registration module 130 (FIG. 1).

Process 800 shifts each of the image blocks of the aligned original document
10 image by its corresponding displacement vector (block 802). Process 800 stitches all the shifted image blocks together to produce a registered original document image (block 804) then terminates.

Shifting all blocks by their assigned displacements and stitching them together (without interpolation) gives a registered original document image.

15 The black pixels are dilated in the registered original using a 5 x 5 structuring element (SE) (that is, each black pixel becomes a 5 x 5 block), forming a mask Im1 containing the "preprinted data".

FIG. 9 is a flowchart illustrating an exemplary authentication process 900 of the authentication module 150 (FIG. 1).

20 Process 900 forms a mask Im1 by dilating the black pixels in the registered original (block 902). Process 900 examines the portion of the final binary copy that is lying outside the mask Im1 (block 904). Process 900 checks if there is any significant region of black pixels in this portion (block 906). Any significant region of black pixels in this portion would indicate that the copy contains additional image
25 material that is not present in the original document image. It is noted that isolated black pixels in this portion are acceptable since they may be caused by noise or dust when the gray-level copy was made from a gray-level original document. If there is any significant region of black pixels in this portion, then process 900 outputs an indicator indicating that the gray-level copy is not authenticated (block 908). Process
30 900 forms a mask M2 by dilating the black pixels in the final binary copy (block 910). Process 900 examines the portion of the registered original document image that is lying outside the mask M2 (block 912). Process 900 checks if there is any significant region of black pixels in this portion (block 914). Any significant region of black

pixels in this portion would indicate that the copy contains less image material than the original document image. If there is any significant region of black pixels in this portion, then process 900 outputs an indicator indicating that the gray-level copy is not authenticated (block 908). Otherwise, process 900 indicates that the gray-level
 5 copy is authenticated (block 916).

FIG. 10 is a flowchart of an exemplary process 1000 of the initial annotation image module 142 (FIG. 1).

Process 1000 forms a mask *lm1* by dilating the black pixels in the registered original (block 1002). Process 1000 binarize the gray-level copy using a very
 10 sensitive threshold to produce a third binary copy (block 1004). In one example, the threshold is set at 220 on a range of 0 to 255, with large values representing brighter gray levels. Process 1000 forms an initial annotation image by using the portion of the third binary copy that lies outside the mask *lm1* (block 1006).

FIG. 11 is a flowchart of the process 1100 of the broken stroke repair module
 15 144 (FIG. 1) as implemented in one embodiment of the present invention.

Process 1100 dilates the mask *lm1* to produce a second mask *lm2* (block 1102), and defines each of connected components located on the initial annotation image in the space between the first mask and the second mask as a broken stroke end (block 1104). Process 1100 forms a list of candidate connections for each
 20 broken stroke end (block 1106). For each broken stroke end, process 1100 computes a connection score for each of its candidate connections (block 1108). For each broken stroke end, process 1100 selects the candidate connection having the highest connection score (block 1110), and accepts the selected candidate connection if its connection score is above an acceptability threshold (block 1112),
 25 then terminates.

Each broken stroke end is usually connected to one other broken stroke end. The exceptional case is when a broken stroke end is itself a connected component in the initial annotation image. It is noted that, in a normal case, a broken stroke end is a part of a connected component in the initial annotation image, but is not itself a
 30 connected component. In this case, process 1100 allows this broken stroke end to be connected to more than one other broken stroke end.

A more detailed description of an implementation of blocks 1106 through 1112 is as follows. First, process 1100 selects candidate pairs to attach one broken stroke

end to another. For example, each broken stroke end x is paired with one other broken stroke end y , and the pair $\{x,y\}$ is put on the list of candidate pairs. In the exceptional case mentioned above, x can be paired again with a different broken stroke end.

5 Process 1100 goes through all the candidate pairs and computes a connection score for each of the candidate pairs. Candidate pairs having the highest connection scores are accepted if these highest connection scores exceed an acceptability threshold.

10 For a pair $\{x,y\}$, process 1100 tests if x has already had some other pair $\{x,z\}$ accepted, and also tests if y has had some pair $\{y,w\}$ accepted. If both x and y have already had candidate pairs accepted, then process 1100 does not accept $\{x,y\}$, except if the $\{x,y\}$ pair was put on the list as part of the exceptional case. In this exceptional case, $\{x,y\}$ pair can be accepted.

15 In one implementation, for each broken stroke end x , up to 5 candidate connections in a center radius are examined. For the normal case, the candidate connection having the highest score is accepted if the connection score exceeds the acceptability threshold. If the broken stroke end x is an exceptional case, then all candidate connections with sufficiently high scores (i.e., above the acceptability threshold) are accepted.

20 The following discussion further clarifies the annotation lifting process. Annotations are often lighter than the printed text, especially if they are in color or pencil. For this reason, it is preferable to rebinarize the grayscale copy using a very sensitive threshold to obtain a third binary version of the gray-level copy (block 1004 in FIG. 10), and take what lies outside mask $Im1$ as the initial annotation image la .

25 Process 1100 improves the appearance of the initial annotation image la by repairing the broken strokes, filling in small gaps where annotations crossed over black pixels in the original document image.

Process 1100 searches for broken strokes not only along preprinted lines, but also at gaps caused by arbitrary interference of the preprinted data and annotations.

30 In one embodiment, process 1100 dilates $Im1$ with a 3×3 structuring element, yielding $Im2$.

An 8-connected component falling in the portion between the two masks, i.e., $Im2 - Im1$, is defined to be a broken stroke end. With one exceptional case, each broken end chooses at most one other broken end to connect to.

FIG. 12 illustrates how the broken stroke ends are defined. In FIG. 12, the cross-hatched portion of the interfering stroke that lies between the two masks $Im1$ and $Im2$ comprises the broken ends.

FIG. 13A illustrates the case where a broken end may be chosen by more than one other broken end. In FIG. 13A, the broken end $c3$ is chosen by both broken ends $c1$ and $c2$.

FIG. 13B illustrates the exceptional case where a broken end is itself a connected component in the initial annotation image la . Such a broken end is allowed to be connect to more than one other broken end. In FIG 13B, broken ends $c2$ and $c3$ are such connected components, so each of them is allowed to connect to two other broken ends. Broken end $c2$ is allowed to connect to $c1$ and $c3$. Broken end $c3$ is allowed to connect to $c2$ and $c4$.

Candidate connections are evaluated according to a line tracing protocol, and are accepted in accordance with a criterion that is based on three principles, which are proximity, consistency and smoothness.

FIG. 14A illustrates the line tracing protocol. To test whether ends $c1$ and $c2$ came from the same stroke, a line is drawn through their centroids $C1$ and $C2$. If $c1$ and $c2$ came from the same annotation stroke, then consistency principle requires that black pixels be present along segment $P1P2$ in the third binary copy and white pixels be present along the segment $P1P2$ in the initial annotation image la . The consistency measure c is set to 1 if this is satisfied, and 0 otherwise.

The smoothness principle requires stroke continuity beyond the points $P1$ and $P2$. The lengths of black pixel runs $l1$ and $l2$ in each direction are computed. The smoothness measure s is defined to be $s = \max \{l1, l2\}$.

The proximity measure p is defined to be distance between $P1$ and $P2$, i.e., $p = |P1 P2|$. Candidate connections having small proximity measure are favored.

The connection score is defined as cs/p and is computed for each candidate connection. Of all possible connections for a given broken end, process 1100 examines the one with the highest connection score, and accepts it if its connection score is sufficiently high. As a practical speed-up, for each broken end, only the five

nearest ends within a certain radius are considered. Thus, the complexity of the line tracing protocol is reduced to linear in the number of broken ends.

FIG. 14B illustrates two broken end cases that are difficult to resolve. The line tracing protocol can robustly resolve such difficult cases.

5 After the candidate connections are accepted, pixels are blackened along accepted connections by dilating the line segment between the two broken stroke ends. The smaller of the two broken stroke ends can be used as the structuring element for this dilation.

10 A system of the present invention has been used to test on a large variety of documents. Annotated copies were generated by multigeneration copying, faxing, and so forth. Different scanners were used for digitizing the original and the annotated copy.

15 FIGs. 15A, 15B, 15C and 15D illustrate one test example. FIG 15A shows the grayscale image of an annotated copy. FIG. 15B shows the initial annotation image. FIG. 15C shows the broken stroke ends. FIG. 15D shows the final annotation image after broken stroke end reconstruction.

20 Various modifications can be made to the embodiment of the present invention as described above. For example, the global alignment process 200 (FIG. 2) works well on document images having small amount of skew and scale variations. To handle larger skew, the global alignment module 110 (FIG. 1) may need to use a preliminary de-skewing process. As another example of modifications, broken stroke repair module 144 (FIG. 1) can use a more sophisticated algorithm than process 1100 for determining stroke width. Smoother gap filling can be achieved using elastic structuring elements.

25 While certain exemplary embodiments have been described in detail and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention. It will thus be recognized that various modifications, such as those mentioned above, may be made to the illustrated and other embodiments of the invention described above, without departing from the broad inventive scope thereof. It will be understood, therefore, that the invention is not limited to the particular embodiments or
30 arrangements disclosed, but is rather intended to cover any changes, adaptations or

modifications which are within the scope and spirit of the invention as defined by the appended claims.

WHAT IS CLAIMED IS

- 1 1. A method for matching an original document image with a copy image, the
2 method comprising the operations of:
 - 3 (a) defining the original document image as an ensemble of blocks, each
4 of the blocks having neighboring blocks;
 - 5 (b) forming a mask for each of the blocks by dilating foreground pixels
6 included in the block;
 - 7 (c) searching for a best match in the copy image, for each of the blocks,
8 using a first search window, by correlating pixels in the mask with pixels in a
9 corresponding portion in the copy image, each of the best matches having a
10 matching score;
 - 11 (d) indicating each of the blocks as "unmatched" when the corresponding
12 matching score is less than a matching threshold; and
 - 13 (e) computing a displacement vector for each of the matched blocks.
- 1 2. The method of Claim 1 further comprising the operation of:
 - 2 (f) binarizing a gray-level copy using a first binarization threshold to
3 produce the copy image.
- 1 3. The method of Claim 2 further comprising the operations of:
 - 2 (g) computing a second binarization threshold based on the matched
3 blocks;
 - 4 (h) binarizing the gray-level copy using the second binarization threshold
5 to produce a second binary copy;
 - 6 (i) assigning to each of the unmatched blocks a displacement vector
7 based on an average of respective displacement vectors of the neighboring
8 blocks of said unmatched block;
 - 9 (j) searching for a new best match in the second binary copy, for each of
10 the blocks, using a second search window, by correlating pixels in the mask
11 with pixels in a corresponding portion in the second binary copy, each of the
12 new best matches having a new matching score;

13 (k) indicating each of the blocks as "unmatched" when the corresponding
 14 new matching score is less than a matching threshold; and
 15 (l) computing a new displacement vector for each of the new matched
 16 blocks.

1 4. The method of Claim 3 further comprising the operation of:

2 (m) repeating operations (i) through (l) until all the blocks are matched.

1 5. The method of Claim 3 wherein the second search window is smaller than the
 2 first search window.

1 6. The method of Claim 1 wherein the blocks are overlapping.

1 7. The method of Claim 1 wherein operation (c) is performed via a hierarchical
 2 search procedure.

1 8. The method of Claim 1 further comprising the operations of:
 2 shifting each of the blocks by the respective displacement vector;
 3 stitching the shifted blocks together to produce a registered original.

1 9. The method of Claim 1 wherein the original document image and the copy
 2 image are in gray-level form.

1 10. The method of Claim 1 wherein the original document image and the copy
 2 image are in binary form.

1 11. The method of Claim 1 wherein the original document image is in binary form
 2 and the copy image is in gray-level form.

1 12. A method for searching for a best match of a block of a document image, the
 2 best match being in a copy image, the method comprising the operations of:

3 (a) defining a search window;

4 (b) defining a first box centered within the search window, the first box
 5 being smaller than the search window;

6 (c) computing matching scores for matches located at points located within
 7 the first box;

- 8 (d) designating the point having the highest matching score as the best
9 match;
10 (e) defining a second box centered around the best match, the second box
11 being smaller than the first box; and
12 (f) repeating operation (c) and (d) with the second box replacing the first
13 box.

1 13. The method of Claim 12 wherein the points in operation (c) are the corners,
2 side midpoints and center of the first box.

1 14. The method of Claim 12 further comprising the operation of:
2 (g) repeating operations (e) and (f) until the points are adjacent to each
3 other.

1 15. The method of Claim 12 further comprising the operations of:
2 (h) defining a second search window around the best match if the best
3 match is located at a boundary of the first search window; and
4 (i) repeating operations (a) through (f).

1 16. A method for producing an annotation image from a copy image and a
2 registered original document image, the method comprising the operations of:
3 (a) forming a first mask by dilating foreground pixels in the registered
4 original document image;
5 (b) forming a first annotation image by using portion of the copy image that
6 lies outside the first mask;
7 (c) forming a second mask by dilating the first mask;
8 (d) defining each of connected components located on the first annotation
9 image in the space between the first mask and the second mask as a broken
10 stroke end; and
11 (e) connecting each of the broken stroke ends to at least one other broken
12 stroke end, the other broken stroke end being chosen in accordance with a
13 line tracing protocol.

1 17. The method of Claim 16 wherein operation (e) comprises the operation of
2 allowing a broken stroke end to be connected to more than one other broken stroke

3 end when said broken stroke end is itself a connected component in the first
4 annotation image.

1 18. The method of Claim 16 wherein operation (e) comprises the operations of:
2 (1) forming a list of candidate connections for a broken stroke end;
3 (2) computing a connection score for each of the candidate connections;
4 and
5 (3) selecting the candidate connection having the highest connection
6 score.

1 19. The method of Claim 18 wherein, in operation (2), the connection score is
2 based on consistency, smoothness and proximity measures.

1 20. The method of Claim 18 wherein, in operation (2), the connection score is
2 computed as cs/p where c represents the consistency measure, s represents the
3 smoothness measure and p represents the proximity measure.

1 21. A system for matching an original document image with a copy image, the
2 system comprising:
3 (a) a partition module for defining the original document image as an
4 ensemble of blocks, each of the blocks including foreground pixels;
5 (b) a mask forming module for forming a mask for each of the blocks by
6 dilating the foreground pixels included in the block;
7 (c) a search module for searching for a best match in the copy image, for
8 each of the blocks, by correlating pixels in the mask with pixels in a
9 corresponding portion in the copy image, the search module assigning to
10 each of the best matches a matching score, and indicating each of the blocks
11 as "unmatched" when the corresponding matching score is less than a
12 matching threshold; and
13 (d) a displacement computing module for computing a displacement vector
14 for each of the matched blocks.

1 22. The system of Claim 21 wherein the partition module defines the blocks such
2 that the blocks are overlapping.

1 23. The system of Claim 21 wherein the search module uses a hierarchical
2 search procedure.

1 24. The system of Claim 21 further comprising:
2 a registration module shifting each of the blocks by the respective
3 displacement vector and stitching the shifted blocks together to produce a
4 registered original.


1 25. A system for producing an annotation image from a copy image and a
2 registered original document image, the system comprising:

3 (a) an initial annotation image module, the initial annotation image module
4 forming a first mask by dilating foreground pixels in the registered original
5 document image, forming a first annotation image by using portion of the copy
6 image that lies outside the first mask, forming a second mask by dilating the
7 first mask, and defining each of connected components located on the first
8 annotation image in the space between the first mask and the second mask
9 as a broken stroke end; and


10 (b) a broken stroke repair module connecting each of the broken stroke
11 ends to at least one other broken stroke end, the other broken stroke end
12 being chosen in accordance with a line tracing protocol.

1 26. The system of Claim 25 wherein the broken stroke repair module allows a
2 broken stroke end to be connected to more than one other broken stroke end when
3 said broken stroke end is itself a connected component in the first annotation image.

1 27. The system of Claim 25 wherein the broken stroke repair module comprises:
2 (1) a connection module for forming a list of candidate connections for a
3 broken stroke end;
4 (2) a score computing module for computing a connection score for each
5 of the candidate connections; and
6 (3) a selecting module for selecting the candidate connection having the
7 highest connection score.



2/14

**FIG. 2**

3/14




FIG. 3

4/14




FIG. 4

FIG. 5






FIG. 6

7/14

**FIG. 8**

8/14






FIG. 9

9/14

**FIG. 10**

10/14

**FIG. 11**

11/14




FIG. 12

12/14




FIG. 13A






FIG. 13B

13/14

**FIG. 14A****FIG. 14B**

14/14




FIG. 15A




FIG. 15B




FIG. 15C




FIG. 15D

Application number - numéro de demande: 2401065

Figures: 7A- 7B - 7C

Pages: _____

Unscannable items
received with this application
(Request original documents in File Prep. Section on the 10th floor)

Documents reçu avec cette demande ne pouvant être balayés
(Commander les documents originaux dans la section de préparation des dossiers au
10^{ème} étage)

