

**METHOD AND APPARATUS FOR CREATING, INDEXING AND
VIEWING ABSTRACTED DOCUMENTS**

BACKGROUND OF THE INVENTION

5

Field of the Invention

The present invention relates to a method
and apparatus for creating a collection of indexed
document images whereby the document images may be
10 retrieved through the index, and to a method and
apparatus for rapidly browsing through document
images by viewing abstract structural views of the
document images rather than the document images
themselves.

15

Description of the Related Art

Recently, as increasingly larger storage
devices have become available, and it has become
possible to store documents not simply as ASCII text
20 but also as a full facsimile image of the document.
More specifically, it is now commonplace to convert
a document into a computer-readable bit map image of



- 2 -

the document and to store the bit map image of the document. Accordingly, whereas ASCII text storage permitted storage and display of only text portions of documents, it is now possible to store a document in computer readable form and to display not only the text but also pictures, line art, graphs, tables and other non-text objects in the document. Likewise, it is also possible to store and display documents such that text attributes, such as size, font, position, etc., are preserved.

Despite these advances, however, it is still difficult to retrieve the document images into computer memory quickly, and then to browse quickly through computer-displayed document images, for example, in a situation where a computer operator retrieves many document images and searches through those document images to find a particular document. These difficulties can be attributed to at least two limitations. First, current limitations on bandwidth of the computer systems limit of the speed at which documents may be retrieved from storage and displayed. For example, at 300 dots-per-inch resolution, an ordinary 8½ by 11 inch black and white document requires approximately 8.4 million bits to store a full document image. Adding halftone (grey levels) or color to the image, or increasing the resolution at which the image is stored, can easily increase storage requirements to many tens of millions of bits. The time required to retrieve those bits from storage and to create and display the resulting image is significant, even with current high speed computing equipment. The time is lengthened even further in situations where the document image is retrieved from storage in a first computer and electronically transmitted, for example, by modem, to a second computer for display on the second computer.

THIS SHEET HAS BEEN
REPLACED




Second, even when a full document image is displayed to an operator, there is ordinarily too much information for an average operator to comprehend quickly. Much of the information displayed to the operator is not relevant to the operator's query and much time is therefore wasted in displaying the non-relevant information. And the presence of such non-relevant information can slow the operator in his attempt to locate and understand document information that is relevant to the query.

Moreover, simply retrieving appropriate documents for presentation to an operator from a large collection of documents can prove difficult because of the large amount of information that must be searched. Conventional document retrieval systems ordinarily rely on the creation of a text index by which text documents may be retrieved. With document images (as opposed to text documents), it has been proposed to subject the document images to optical character recognition processing ("OCR processing") and to index the resulting text. Systems such as those proposed in U.S. Patent 5,109,439 to Froessl suggest that is only necessary to OCR-process specific areas of the document to simplify the indexing process, but it has nevertheless heretofore proved difficult to create an adequate index for retrieval of document images.

SUMMARY OF THE INVENTION

It is an object of the present invention to address the foregoing difficulties.

In one aspect, the invention provides a system for rapidly browsing through document images by displaying not the full document image, but rather only an abstract structural view of the



document image. More specifically, a structural
view of a document image includes labelled objects
such as "title", "text", "table", "line art", "half-
tone" and similar objects that are ordinarily found
5 in documents, and those objects are displayed for
presentation to the operator in an abstract
structured layout that is representative of the full
document image. The level of abstraction and how
the document is rendered on the display may be
10 defined by the operator. Thus, it is possible for
the operator to specify a high level of abstraction
where only labelled objects are displayed in the
appropriate layout; likewise, it is also possible
for the operator to specify a lower level of
15 abstraction where the original document image for
only specific types of objects, such as for titles,
is displayed. It is also possible for the operator
to specify that text type objects should be
displayed as ASCII equipment characters rather than
20 the original image of the text. By displaying only
abstract structural views of document images rather
than the full document image, and by permitting the
operator to specify the level of abstraction, it is
possible to retrieve and display document images
25 quickly and to present the information to the
operator in a quickly comprehensible format.

According to this aspect, the invention
provides a method and apparatus for storing document
images in which a document image is obtained, for
30 example, by scanning an original document, the
document image is processed to determine the
structure of the document, and the structure of the
document is stored together with the document image.
Preferably, the structure of the document is
35 determined by conventional block selection
techniques which utilize a rule-based knowledge
system for identifying specific areas in a document



- 5 -

and for determining the content of the image within those areas such that the document image is decomposed into a general set of objects. Suitable block selection techniques are described in co-
5 pending application Serial No. 07/073,012 filed April 24, 1992 entitled "Method and Apparatus for Character Recognition".

According to this aspect, the invention further provides a method for displaying document
10 images in which, given an image of a document and its corresponding structure, a representation of the structure of the document is displayed and selected structural areas of the display are selectively replaced with corresponding areas from the full
15 document image. The abstraction level at which the structure is displayed may be designated, whereby the structural representation may be displayed with mixed areas of structure and full document image (or ASCII text equivalent) in accordance with the
20 designated abstraction level. If desired, the display may be either a visual display or a verbal display in which the structure and content of the document is enunciated through text-to-speech techniques which convert the structural and
25 contextual information into spoken words for aural communication.

In a further aspect of the invention, a retrieval index by which images of documents may be retrieved is created. According to this aspect a
30 document is scanned to form a document image, and the document image is processed, for example, by block selection techniques, to identify areas on the document of a first type such as title areas, and areas on the document of other types. The document
35 image in the first type areas are converted into text, for example, by optical character recognition ("OCR") techniques. The converted text is indexed



- 6 -

so as to form a retrieval index, and the document image is stored together with the indexed text such that the stored document index may be retrieved with reference to the indexed text. Advantageously, the invention creates a retrieval index only from first-type areas of the document, such as from title areas of the document or from the first paragraph of text areas of the document.

This brief summary has been provided so that the nature of the invention may be understood quickly. A more complete understanding of the invention can be obtained by reference to the following detailed description of the preferred embodiment thereof in connection with the attached drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a perspective view of the outward appearance of an apparatus according to the invention, and Figure 2 is a block diagram of the Figure 1 apparatus.

Figure 3 is a full diagram showing the overall operation of the Figure 1 apparatus.

Figure 4 is a full diagram showing inputting and indexing documents according to the invention.

Figure 5 is a view showing a representative input document.

Figure 6 is a view for explaining how various blocks identified within original documents are stored.

Figure 7 is a flow diagram for explaining document display according to the invention.

Figures 8 to 10 are views showing representative displays of an abstracted document,



- 7 -

in which Figures 9 and 10 are views showing expansion of selected areas of the abstracted document.

5


DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Figures 1 and 2 show an apparatus according to the present invention.

10 As shown in these figures, reference numeral 10 designates personal computing equipment such as an IBM PC or PC-compatible computer. Computing equipment includes a CPU 11 such as an 80386 processor which executes stored program
15 instructions such as operator selected applications programs that are stored in RAM 12 or specialized functions such as start-up programs or BIOS which are stored in ROM 14. Computing equipment 10
20 further includes a local area network interface 15 which provides interface to a local area network 16 whereby the computing equipment 10 can access files such as document files on a remote file server or send files for remote printing or have remote
25 machines access document images on equipment 10 or otherwise interact with a local area network in accordance with known techniques such as by file exchange or by sending or receiving electronic mail.

Computing equipment 10 further includes a monitor 17 for displaying graphic images and a
30 keyboard/mouse 19 for allowing operator designation of areas on monitor 17 and inputting information.


Mass storage memory 20, such as a fixed disk or a floppy disk drive, is connected for access by CPU 11. Mass storage 20 typically includes
35 stored program instruction sequences such as an instruction sequence for scanning, indexing, retrieving and displaying documents according to the



- 8 -

invention, as well as other stored program instruction sequences for executing application programs such as word processing application programs, optical character recognition programs, block selection applications programs, spreadsheet application programs, and other information and data processing programs. Mass storage memory 20 further includes document index tables which contain index information by which documents may be retrieved as well as bit map images of documents, document structures, and ASCII text for text areas of the documents. Other data may be stored on mass storage memory 20 as desired by the operator.

A modem 21, a facsimile interface 22, and a voice telephone interface 24 are provided so that CPU can interface to an ordinary telephone line 25. Each of the modem 21, facsimile interface 22, and voice telephone interface 24 are given access to the telephone line 25 via a telephone line switch 26 which is activated under control by CPU 11 so as to connect telephone line 25 to one of the modem 21, the facsimile 22, or the voice telephone interface 24, as appropriate to the data being sent and received on the telephone line. Thus, CPU 11 can send and receive binary data such as ASCII text files or document images files via modem 21 and it can be controlled by a remote computer via modem 21, it can send and receive facsimile messages via facsimile interface 22, and it can interact on an ordinary voice telephone line via voice telephone interface 24. In this regard, voice telephone interface 24 is provided with a DTMF decoder 24A so as to decode tones on the voice telephone line 25 which correspond to operator depressions of a telephone keypad. In accordance with stored program instruction sequences in mass storage memory 20, the decoded tones are interpreted by CPU 11 into



operator commands, and those operator commands are executed so as to take predesignated actions in accordance with operator depressions of the telephone keypad.

5 A conventional text-to-speech convertor 27
is connector to the CPU 11. The text-to-speech
convertor 27 interprets text strings that are sent
to it and converts those text strings to audio
speech information. The text-to-speech convertor 27
10 provides audio speech information either to a
speaker 27 for enunciation to a local computer
operator, or provides audio speech information to
the voice telephone interface 24 for enunciation
over ordinary voice telephone lines.

15 MIDI ("Musical Instrument Digital
Interface") synthesizer 30 is also connected to
CPU 11 and interprets MIDI music commands from
CPU 11 so as to convert those MIDI music commands to
audio wave forms. The audio wave forms are, in
20 turn, played out over speaker 28 or provided to
voice telephone interface 24 for play out over
ordinary voice telephone lines.

Scanner 31 operates to scan original
documents printed on a sheet of paper, and to
25 convert the information of those original documents
into a bit-by-bit computer readable representation
of that document. Scanner 31 may be a simple black
and white scanner, but more preferably scanner 31
includes at least half-tone (grey scale) processing
30 capabilities and/or color processing capabilities.

Printer 32 is provided to form images of
documents under the control of CPU 11. Printer 32
may be an ordinary black and white printer, but,
more preferably, printer 32 includes half-tone
35 and/or color capabilities.

A CD ROM 34, such as an optical disk, is
connected for access by CPU 11. CD ROM 34 operates



- 10 -

to supplement the storage in mass storage memory 20 and contains additional information concerning document indexes and document images and document structure. It is also possible to provide a write-once-read-many ("WORM") optical device or an ordinary read/write optical device so as to further supplement the storage capabilities. In addition, the local area network 16, CPU 11 can access document indexes and document images and document structure stored at remote file server locations, and via modem 21, CPU 11 can access document indexes and document images stored at centralized data base locations over ordinary voice telephone lines.

Figure 3 is a flow diagram showing overall operation of the Figures 1 and 2 apparatus. The process steps shown in Figure 3 are executed by CPU 11 in accordance with stored program instruction sequences stored in mass storage 20 and executed out of RAM 14.

In step S301, CPU 11 displays a main selection menu on monitor 17 and permits operator selection of one of the menu items. In step S302, if the operator has selected to input documents, then flow advances to step S303 in which document image processing is executed, after which flow returns to step S301. If the operator does not select document input but instead selects to retrieve documents (step S304) then flow advances to step S305 in which document retrieval processing is executed, after which flow returns to step S301. If the operator does not select document retrieval but instead selects to display documents (step S306), then flow advances to step S307 in which document display processing is executed, after which flow returns to step S301. Each of these processes is discussed in additional detail below.



- 11 -

Although the flow processing illustrated in Figure 3 appears to be executed serially and sequentially, it is to be understood that the processes may be executed at significantly different times, in different order, and indeed may be executed by different machines. For example, it is possible for a first machine to input documents and to store those documents for retrieval and display by a different machine. Likewise, it is possible for one machine to retrieve documents and transmit them to another machine, for example, via modem 21, for display on that other machine.

Figure 4 is a detailed flow diagram illustrating document input processing in accordance with step S303 of Figure 3.

In step S401 CPU 11 causes scanner 31 to scan an original document fed into scanner 31 so as to convert the original document into a computer readable bit map image of that document. The bit map image may be temporarily stored in RAM 14 for further processing.

In step S402, CPU 11 derives structural information concerning the layout of objects within the document. More specifically, CPU 11 subjects the bit map document image to block selection techniques such as those described in co-pending U.S. Application, Serial Number 07/873,012 so as to decompose the document image into a general set of objects. Using rule-based decisions, such block selection techniques are able to identify related regions or areas within the document, and are able to determine the type of images within those regions. For example, using the block selection techniques described in the aforementioned U.S. Application 07/873,012, CPU 11 processes the scanned-in bit map image stored in RAM-14 to identify related regions in the image. Then, CPU 11



identifies the type of region, for example, whether a region is a title region, a text region, a paragraph region, a table region, a line art image region such as a graph, a half-tone image region, a color image region, etc.

Figure 5 illustrates regions identified on a document image in accordance with processing according to the process steps in Step 402. Figure 5 is a representative document image consisting of page 40 from a magazine article. As shown in Figure 5, page 40 includes title regions such as regions 41, 42 and 47, as well as text regions such as regions 43, 44 and 48. Likewise, page 40 includes line art image regions 45, table regions 49 and half-tone object regions 46. Other types of regions are possible. In accordance with processing in Step S402, CPU 11 identifies the related regions in the document and determines the type of image in each of the regions. In Figure 5, each region has been indicated diagrammatically with a circumscribing rectangle.

Reverting to Figure 4, flow then advances to Step S403 in which CPU 11 identifies the document type. More specifically, many documents have relatively standard structures; for instance, a business letter virtually always includes a corporate letterhead region (which will be identified in Step S402 as a line art image area), a date and an addressee region (which will be identified in Step S402 as separate but adjacent text regions), salutation, body and closing regions (which will be identified as text regions in Step S402), and a signature (which will be identified as a line art image region). CPU 11 stores templates of these standard document types, and in Step S403 compares the structural information derived in Step S402 with those stored templates. If a match is



- 13 -

found between the stored template and the structural information derived in Step S402, then CPU 11 assigns the associated document type to the document image stored in RAM 14.

5 In Step S404, CPU 11 correlates related regions within the document. More specifically, there are often regions within a document that should be considered together, such as the text caption for a figure or the footnote for a text
10 block. For example, referring to Figure 5, text region 48 is related to title 47, and table region 49 is related to title 50. Using rule-based decision processes, CPU 11 correlates related regions.

15 In Step S405, CPU 11 subjects designated document regions to OCR processing in preparation for creating a retrieval index for the document. More specifically, so as to be able to retrieve the document image based on a word-based retrieval
20 system, CPU 11 creates a retrieval index from text regions within the document. It is possible to subject all text regions identified in Step S402 to OCR processing and to form the index based on the resulting text, but more preferably CPU 11 offers
25 the operator a selectable level of indexes. More specifically, CPU 11 offers the operator the option of selecting the index level in accordance with the following levels:

- Level 1: Title Regions Only
- 30 Level 2: First Line Of Text Regions
- Level 3: Text Captions Related To Graphics
- Level 4: Full Text

 If the operator selects Level 1 ("Titles Only"), then CPU 11 subjects only title regions to
35 OCR processing. The resulting text from the title regions is provided for a retrieval index to allow the document to be retrieved when words in the title




- 14 -

match a specified search query. Likewise, if the operator specifies Index Level 2 ("First Line of Text Regions"), then CPU 11 subjects both title regions as well as only the first line of each text regions to OCR processing, and creates a retrieval index based on the resulting text. If Index Level 3 is selected ("Text Captions Related To Graphics"), then CPU 11 subjects the items in Levels 1 and 2 to OCR processing as well as text regions related to graphics as specified in the correlations obtained in step S404. Finally, if the operator specifies Index Level 4 ("Full Text"), then both title regions and full text regions wherever located are subjected to OCR processing and the resulting text provided to form the retrieval index. It is possible to include other indexing levels.

In step S406, the ASCII text characters obtained from OCR processing are indexed and the document retrieval index is updated appropriately. In addition to indexing the ASCII text words, a flag is also indexed which indicates the type of object from which the word was obtained, that is, title, text, caption, etc. The flag may be used as part of the retrieval process so as to retrieve documents based on the occurrence of words in particular regions of documents.

In Step S407, CPU 11 stores the structural information that was derived in Step S402 and that was correlated in Step S406. Figure 6 is a representative view of how the structural information is stored. As shown in Figure 6, for each document the structural information includes a document identifier 51 which is the same identifier as that assigned to the full document image and by which it is possible to retrieve the full document image. In area 52, the document type derived in Step S402 is stored. At 53, all of the structural



- 15 -

information for the document, and its layout within the document, is stored. As shown at 53, for each region identified in Step S402, CPU 11 stores a region identifier, a region type, rectangular coordinates that specify the position of the region on the page (here the upper left corner ("ULC") coordinates and the lower right corner ("LRC") coordinates), and all related regions identified in Step S406. In Figure 6, region 1 corresponds to region 41 in Figure 5 and, as shown in Figure 6, includes a "title" type, upper left hand coordinates of $\langle 0,0 \rangle$, lower right hand coordinates of $\langle 5,40 \rangle$, and no related regions. The remaining regions illustrated in Figure 6 follow in like fashion.

Region attributes 54 may also be stored, for example, font selection and size as illustrated in Figure 6. Other attributes may also be stored, such as number of text lines, number of paragraphs, etc.

In Step S408, the document image is stored, preferably with OCR-processed document text. The document text is stored in the event that a quick display of ASCII text characters is desired rather than the slower display of images of text regions. Accordingly, at the end of Step S408, CPU 11 has completed document input processing and has stored a text-searchable index for the document, a full image for the document, the structure of the document, and, if desired, OCR-processed text of text regions of the document.

Reverting to Figure 3, if an operator selects document retrieval (Step S304), then flow advances to Step S305 for document retrieval processing. Document retrieval processing used in the preferred embodiment of the invention is a word-based retrieval and search system based on an operator-specified search query. Such a query may



be specified by a logical sequence that must be found in each document (a "Boolean" search), or may be specified by a natural language based query system. The query may be made directly by a local operator, or remotely by an operator with voice or DTMF phone access, or from commands issued by a remote computer via modem or LAN.

The operator may also specify that documents are retrieved based on the occurrence of words in particular document regions, such as title regions or captions. Documents matching the query are retrieved and are assigned scores based on the degree that the document matches the query. Higher scores may be assigned for occurrences of words in particular regions of documents, the score being scaled the same as the above indexing level selection. Thus, for example, a higher score may be assigned for occurrences of words in title regions than for occurrences of words in text regions. When plural documents are retrieved, they are displayed in order of their score.

When an operator selects document display processing (Step S306), then flow advances to Step S307 for CPU 11 to execute document display processing. Figure 7 is a detail flow diagram illustrating document display processing.

In Step S701, CPU 11 allows the operator to set display options by which retrieved documents will be displayed. Specifically, the operator selects one of two modes by which retrieved documents will be displayed: full document image mode or browse mode. If the full document image mode is selected, then full document images of retrieved documents will be displayed. If the browse mode is selected, then only an abstract structural view of each document image is displayed, thereby permitting an operator to browse rapidly



- 17 -

through all retrieved documents. In the browse mode, three further parameters may be set by the operator. First, the operator may select whether to display ASCII text characters of OCR-processed text regions or document images of text regions. Second, the operator may select the abstraction level at which the structural view is displayed. Nine abstraction levels are possible in this embodiment of the invention, though other arrangements are possible:

- Level 1: Structure Only
- Level 2: Titles
- Level 3: Retrieval Parameters
- Level 4: First Line of Text
- Level 5: Full Text
- Level 6: Tables
- Level 7: Line Art
- Level 8: Half-tone
- Level 9: Color

In accordance with the selected abstraction level, CPU 11 will display structural views of the document mixed with regions of full document images. If abstraction Level 1 is selected, then only structure is displayed. But if abstraction Level 2 or above is selected, then structural information is displayed mixed with regions of full document images. For example, at abstraction Level 2, structural information is displayed mixed with images (or ASCII text, if selected) of title regions of the document. The needed image regions are obtained from the full document image in accordance with the ULC and LRC coordinates stored with the structure. If abstraction Level 3 is selected, then structural information is displayed mixed with images of titles (or ASCII text, if selected) and the retrieval parameters by which the document was retrieved. At



- 18 -

abstraction Level 4, structural information is displayed mixed with images of title regions, retrieval parameters, and the first line of text images in each text block. As before, if ASCII text display is selected, then ASCII text is displayed rather than images of text. Higher abstraction levels cumulatively display more and more of the full document image in accordance with the selected abstraction level.

In addition to the abstraction level parameter, in the browse mode the operator can also select whether to allow image enhancement manually or automatically. Specifically, in manual image enhancement, when the structure of a retrieved document is displayed, the operator may target a particular region of the document, for example with mouse 19, and request for CPU 11 to replace the displayed region with the full document image. If automatic enhancement is selected, then CPU 11 automatically replaces the structural information that it is displaying with full document images. Automatic enhancement operates during quiescent periods of CPU operation, that is, when the CPU has sufficient processing time so as to enhance the image. Enhancement progresses in the order of the abstraction levels given above. Thus, during quiescent processing periods, CPU begins to replace, first, title regions with full images of the title regions, second, retrieval parameter regions with full images of the retrieval parameters, third, the first lines of the text regions with full images of the first lines of text, etc.

After display options have been set in Step S701, flow advances to Step S702 in which CPU 11 determines whether the full image mode has been selected. If the full image mode has been selected,



then flow branches to Step S703 in which full images of the full images are displayed.

On the other hand, if the browse mode has been selected, then flow advances to Step S704 in which the structure of the document is displayed at the selected abstraction level. A representative display of structural information is shown in Figure 8. In Figure 8, the operator has selected the browse mode, and has selected ASCII text display at abstraction Level 2, that is "Titles". Thus, as shown in Figure 8, an abstracted structural view of the document is presented to the operator, with ASCII text equivalents of title regions. The remaining regions are simply labeled with the region type, as specified in the region type area of Figure 6. Labelling may be by color coding of the region, or, as shown in Figure 8, by applying a text label. Each of the regions is circumscribed by a rectangular shape whose coordinates are specified by the "ULC" and "LRC" coordinates of Figure 6. ASCII text equivalents of title regions are displayed at the corresponding ULC and LRC coordinates. Thus, the operator is presented with an abstract structural view of the document with a selectable abstraction level. This permits the operator to browse quickly through retrieved documents so as to find particular documents of interest.

Flow then advances to step S705 in which CPU 11 determines if the manual enhancement parameter has been selected. If the manual enhancement parameter has not been selected, then flow branches to step S706 in which CPU 11 automatically replaces structurally displayed regions of the document with full images (or ASCII text, if selected) of those regions during quiescent operating times of the CPU. Thus, assuming that sufficient quiescent operating time is available, if



the automatic enhancement mode has been selected, then the full document will eventually be presented to the operator as the CPU replaces the structurally represented regions of that document. At any time
5 during this process the operator may select for the next retrieved document to be displayed (Step S710) whereupon flow returns to Step S702 so as to display the next retrieved document.

If the manual enhancement mode has been
10 selected, then flow advances to Step S707 in which CPU 11 determines whether a document region has been targeted by the operator. If no document region has been targeted by the operator, then CPU 11 does nothing, or more preferably, begins to replace
15 structural regions of the document with full document images as shown at Step S708. Again, this process may be interrupted at any time by operator selection of the next document (Step S710).

On the other hand, if in Step S707 the
20 operator targets a document region for replacement, then flow advances to Step S709 in which CPU 11 retrieves the corresponding document image form the document image file and replaces the structurally represented region of the displayed document with
25 the full document image. Thus, as shown in Figure 9, the operator has targeted area 54 which is a text region. CPU 11 thereupon retrieves the document image corresponding to the rectangular area circumscribed by the ULC and LRC coordinates, or the
30 ASCII text equivalent if that option has been selected, and replaces the area in Figure 8 with the corresponding document image as shown in Figure 9. Flow then returns to Step S707 to await further designations by the operator.

35 In Figure 9, the ASCII text equivalent for region 54 does not fit into the allotted area. In this situation, CPU 11 only displays the amount of



- 21 -

text that will fit, but, so as to allow the operator to view all the text, CPU 11 permits scrolling within the area.

5 As further shown in Figure 9, the operator has requested for the first line of text from region 44 to be displayed, and the ASCII text equivalent of that text is displayed at 55. Of course, if image mode was selected, then the image of the first line of text would have been displayed.

10 In Step S709, if an operator targets a document region which has a related region, such as region 42 or 43, then not only does CPU 11 replace the targeted region, but CPU 11 also replaces the related region as well. Accordingly, if the
15 operator targets region 42, which is a title region, CPU 11 replaces title region 42 with the full image (or text) of region 42, and also replaces text region 43 with the full image (or text) of text region 43. By virtue of this feature, the operator
20 is provided with quick access to information that is desired without requiring the operator to execute additional keystrokes, etc.

Figure 10 shows a situation in which the operator has selected for certain regions to be
25 displayed as images rather than as ASCII text. More particularly, the operator has requested for full images of regions 46 and 48, and those regions have been provided as shown at 56 and 57, respectively.



WHAT IS CLAIMED IS:

1. A method for creating a retrieval index by which images of documents may be retrieved, the method comprising the steps of:

processing the document image to identify areas on the document of a first type and areas on the document of other types;

converting the document image in the first type area into text;

indexing the text; and

storing the document image such that the stored document image may be retrieved by reference to the indexed text.

2. A method according to Claim 1, further comprising the step of storing the identity of the regions identified on the document.

3. A method according to Claim 2, further comprising of the step of retrieving the stored identities and displaying a representation of the document in accordance with the stored identities.

4. A method according to Claim 3, further comprising the step of progressive replacement of identified regions in the representation of the document with full images of the document.

5. A method according to Claim 4, wherein the progressive replacement is automatic.

6. A method according to Claim 4, wherein the progressive replacement is manual in accordance with operator designation.



7. A method according to Claim 3, wherein documents are displayed by score.

8. A method according to Claim 1, further comprising the step of selectably designating the first type of region for indexing.

9. A method according to Claim 8, wherein the first type of region is selected from the group consisting of title type regions and text type regions.

10. Apparatus for creating a retrieval index by which images of documents may be retrieved, said apparatus comprising:

- a scanner for scanning an original document and for outputting a bit map document image;

- a first memory region for storing a document image and a retrieval index;

- a second memory region for storing process steps; and

- a processor for executing the process steps stored in said second memory;

wherein said second memory includes process steps to (a) receive the bit map image scanned by said scanner, (b) process the bit map image to identify areas in the document of a first type and areas in the document of other types, (c) convert of the document image in the first type region into text, (d) index the text and update the index stored in said first memory with the indexed text, and (e) store the document image in said first memory such that the stored document image may be retrieved by reference to the indexed text.



This Sheet Has Been Replaced



- 24 -

11. Apparatus according to Claim 10, wherein said second memory includes process steps to store the identity of the regions identified in the bit map document in the bit map image in said first memory.

12. Apparatus according to Claim 11, further comprising a display, and wherein said second memory includes process steps to display a representations of retrieved documents on said display in accordance with the identity of the document regions.

13. Apparatus according to Claim 12, wherein said second memory includes process steps to progressively replace regions in the displayed document representation with corresponding regions from the full document image.

14. Apparatus according to Claim 13, wherein the progressive replacement is automatic.

15. Apparatus according to Claim 14, further including a designation means, and wherein the replacement is manual in accordance with designation by said designation means.

16. A method according to Claim 12, wherein documents are displayed by score.


17. Apparatus according to Claim 10, wherein the first type is selectably designatable.

18. Apparatus according to Claim 17, wherein the selectable designations over the first type are selected from the group consisting of title type regions and text type regions.

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 25 -

19. A document storage method comprising the steps of:

scanning a document to form a document image;

processing the document image to determine the structure of the document; and

storing the structure of the document in correspondence with the document image.

20. A method according to Claim 19, further comprising the step of indexing selectable types of document structure.

21. A method according to Claim 19, further comprising the step of identifying a document type based on the determination of the document structure, and storing of the identified document type.

22. A method for displaying documents comprising the steps of:

providing an image of a document and a corresponding structure for the document;

displaying a representation of the structure of the document;


replacing selected regions of the displayed document representation with corresponding document regions.

23. A method according to Claim 22, further comprising the step of designating an abstraction level for displaying the representation of the document, and wherein in said displaying step the representation is displayed with mixed areas of structure and document regions in accordance with the designated abstraction level.

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 26 -

24. A method according to Claim 22, wherein in said providing step a plurality of document images and corresponding document structures are provided, and further comprising the step of retrieving one of the plurality of document images and corresponding document structures for display.

25. A method according to Claim 24, wherein the documents are displayed by score.

26. A method according to Claim 24, wherein in said displaying step, a document structure is displayed mixed with retrieval parameters.

27. A method according to Claim 22, further comprising the step of labelling the structural regions in the displayed representation of the document.

28. A method according to Claim 22, wherein said displaying step is a verbal enunciation of the document structure and/or content through text-to-speech conversion.

29. A method according to Claim 22, further comprising the step of identifying document type based on the document structure.

30. A method according to Claim 22, further comprising the step of identifying related regions in the document, and wherein in said replacing step, when a selected region of the displayed document representation is replaced, related regions are also replaced.

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 27 -

31. A method according to Claim 22, wherein in said replacing step, text type document regions are replaced with images of document regions or text equivalents in accordance with a manual selection.

32. Apparatus for displaying a document comprising:

means for providing an image of a document and a corresponding structure for the document;

means for displaying a representation of the structure of the document;

means for replacing selected regions of the displayed document representation with corresponding document regions.

33. An apparatus according to Claim 32, further comprising designating means for designating an abstraction level for displaying the representation of the document, wherein said representation is displayed with mixed areas of structure and document regions in accordance with the designated abstraction level.

34. An apparatus according to Claim 32, wherein a plurality of document images and corresponding document structures are provided, and further comprising retrieving means for retrieving one of the plurality of document images and corresponding document structures for display.


35. An apparatus according to Claim 34, wherein the documents are displayed by score.

36. An apparatus according to Claim 34, wherein said displaying means displays a document structure mixed with retrieval parameters.

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 28 -

37. An apparatus according to Claim 32, further comprising labelling means for labelling the structural regions in the displayed representation of the document.

38. An apparatus according to Claim 32, wherein said displaying means is a verbal enunciation of the document structure and/or content through text-to-speech conversion.

39. An apparatus according to Claim 32, further comprising identifying means for identifying the document type based on the document structure.

40. An apparatus according to Claim 32, further comprising identifying means for identifying related regions in the document, and wherein when said replacing means replaces a selected region of the displayed document representation, related regions are also replaced.

41. An apparatus according to Claim 32, wherein text type document regions are replaced with images of document regions or text equivalents in accordance with a manual selection.

42. A document display apparatus comprising:

a display;

a first memory region from which an image of a document and corresponding structure for the document may be retrieved;

a second memory region for storing processing steps;

a processor for executing the process steps stored in said second memory;

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 29 -

wherein said second memory includes process steps to (a) display on said display a representation of the structure of a retrieved document, and (b) replace selected regions of the displayed document representation with the corresponding document regions.

43. Apparatus according to Claim 42, further comprising input means for inputting operator designations, and wherein said second memory includes process steps to permit designation of an abstraction level at which document representations are displayed, wherein document representations are displayed with mixed areas of structure and document in accordance with the designated abstraction level.

44. Apparatus according to Claim 42, wherein said first memory includes plural document images and plural corresponding structures, and wherein said second memory includes process steps to retrieve one of the plurality of stored document images and corresponding document structures for display.

45. An apparatus according to Claim 44, wherein the documents are displayed by score.


46. Apparatus according to Claim 44, wherein document structure is displayed mixed with retrieval parameters.

47. Apparatus according to Claim 42, wherein said processor labels the displayed document representation in accordance with the structure stored in said first memory.

THIS SHEET HAS BEEN
REPLACED



This Sheet Has Been Replaced



- 30 -

48. Apparatus according to Claim 42, wherein said processor enunciates the document structure in accordance with text-to-speech conversion.

49. Apparatus according to Claim 42, wherein said processor identifies document type based on the stored document structure.

50. Apparatus according to Claim 42, wherein the stored document structure correlates related regions in the document, and wherein when said processor replaces a selected region of the displayed document representation with corresponding regions of document image, said processor also replaces related regions.

51. An apparatus according to Claim 42, wherein text type document regions are replaced with images of document regions or text equivalents in accordance with a manual selection.

THIS SHEET HAS BEEN
REPLACED