AN INTEGRATED APPROACH TO DOCUMENT DECOMPOSITION AND STRUCTURAL ANALYSIS

Debashish Niyogi and Sargur N. Srihari

Center of Excellence for Document Analysis and Recognition
State University of New York at Buffalo
UB Commons, 520 Lee Entrance, Suite 202
Amherst, NY 14228-2567

Abstract

A document image is a visual representation of a paper document, such as a journal article page, a cover page of facsimile transmission, office correspondence, an application form, *etc.* Document image understanding as a research endeavor consists of developing processes for taking a document through various representations: from scanned image to semantic representation. This paper describes document decomposition and structural analysis, which constitutes one of the major processes involved in document image understanding. The current state-of-the-art and future directions in the areas of document segmentation, layout analysis, and logical block grouping are indicated. A system that performs decomposition and structural analysis (including logical grouping and read-order determination) on complex multi-articled documents is presented. This system uses bottom-up segmentation techniques to identify the block structure of a document, and layout rules to classify and group these blocks into logical units that represent meaningful subdivisions of the document. Experimental results showing the efficiency of this approach are presented and discussed.

Contents

List of Figures

1 Introduction

The need to process documents on paper by computer has led to an area of research that may be referred to as *document image understanding* [DIU]. The goal of a DIU system is to convert a raster image representation of a document (*e.g.,* a paper document scanned by a flatbed document scanner) into an appropriate symbolic form [1]. DIU as a research endeavor consists of studying all processes involved in taking a document through various representations: from a scanned or facsimile multi-page document to high-level semantic descriptions of the document. Thus it involves many sub-disciplines of computer science including image processing, pattern recognition, natural language processing, artificial intelligence and database systems.

The symbolic representation desired as output of a DIU system can take one of several forms: an editable description, a representation from which the document can be (exactly) reconstructed, a semantic description useful for document sorting/filing, *etc.* Representation schema that are useful for editing and exact reproduction are standards for electronic document description. An example of a goal-oriented representation useful for document storage and processing is the assignment of a ZIP+4 Code to a document (such as a postal envelope) from an address contained in it. The task involves looking up postal directories for matching against city, state, ZIP Code and the street address.

Developing a DIU system with performance comparable to that achieved by a human expert is still decades from realization [2]. The state-of-the-art in DIU can be subdivided into the following areas:

1. *System architecture* – The complexity of the DIU task leads to modularization into manageable processes. Due to interdependency of processes, issues of how to maintain communication and integrate results from each process arise.

2. *Decomposition and Structural Analysis* – Documents consist of text (machine-printed and handwritten), line drawings, tables, maps, half-tone pictures, icons, etc. It is necessary to *decompose* a document into its component parts in order to process these individual components. Their *structural analysis*, in terms of spatial relationships and logical ordering, is necessary to invoke modules in appropriate order and to integrate the results of the appropriate modules.

3. *Text recognition and interpretation* – It is necessary to recognize words of text, often using lexicons and

higher level linguistic and statistical context. The necessity for contextual analysis arises from the fact that it is often impossible to recognize characters and words in isolation, particularly with handwriting and degraded print.

4. *Tables, graphics and halftone recognition* – Specialized subsystems are necessary for processing a variety of non-text or mixed entities, such as recognizing tabular data, converting graphical drawings into vector representation, and extracting objects from half-tone photographs.

5. *Databases and system performance evaluation* – Methods for determining data sets on which evaluation is based and the metrics for reporting performance.

In this paper, we will describe an integrated system for document decomposition and structural analysis (the second of the above areas), which is a critical process within DIU.

2 Functional Characteristics of a DIU System

The input to a typical DIU system is a high-resolution gray-scale image (*e.g.,* 300ppi). Using gray-scale imagery increases image analysis capability, *e.g.*, gray-scale character recognition [3] and analysis of half-tone photographs. The output of the system is a description of the contents of the document components. An editable description should contain the following entries: (i) component identities and locations on the document, *e.g.,* text, graphics, half-tones, *etc.*, (ii) spatial relationships between components, (iii) layout attributes, *e.g.,* component size, number of lines in text blocks, *etc.*, and (iv) logical grouping of components.

The level of document understanding is determined by the level of representation that the DIU system can derive from a document image. Three types of information can be derived from a document: (i) *layout (geometric) structure* is a physical description of document regions, *i.e.,* size, location, spatial relationships between regions, (ii) *logical structure* is a grouping of layout components based on *human* interpretation of the content of the components and the spatial constraints between components, and (iii) *content interpretation* contains coded data of a component which can be used to derive logical structure or can be stored for later access. A system whose output contains only layout structure is a *document layout analysis system* and a system whose output contains all three types of information is a *DIU system.* From this perspective, layout analysis is an intermediate step of DIU [4].

3 Document Decomposition and Structural Analysis

A *document* image is a visual representation of a printed page such as a journal article page, a facsimile cover page, a technical document, an office letter, *etc.* Typically, it consists of blocks of text, *i.e.*, letters, words, and sentences that are interspersed with tables and figures. The figures can be symbolic icons, gray-level images, line drawings, or maps. A digital document image is a two-dimensional representation of a document obtained by optically scanning and digitizing a hardcopy document. It may also be an electronic version that was created for publishing or drawing applications available for computers.

Methods of deriving the blocks can take advantage of the fact that the structural elements of a document are generally laid down in rectangular blocks aligned parallel to the horizontal and vertical axes of the page. The methods can also use several types of knowledge including *visual, spatial,* and *linguistic*. Visual knowledge is needed to determine objects from the background. Labeling blocks involves the use of spatial knowledge, *e.g.*, the layout of a typical document. Determining the font and identity of characters also involves spatial knowledge. Reading words in degraded text is a process that involves spatial as well as linguistic knowledge, *e.g.*, a lexicon of acceptable words. Determining the role of a block of text, *e.g.*, "is this a title?", is a process requiring spatial, syntactic as well as semantic knowledge. Considerable interaction among different types of knowledge is necessary. For instance, assigning a role to a textual region may require not only knowledge of the spatial layout, but also an analysis of its textual syntax and semantics, and an interpretation of neighboring pictorial regions.

The document decomposition and structural analysis task can be divided into three phases [1]. Phase 1 consists of *block segmentation* where the document is decomposed into several rectangular blocks. Each block is a homogeneous entity containing one of the following: text of a uniform font, a picture, a diagram, or a table. The result of phase 1 is a set of blocks with the relevant properties. A textual block is associated with its font type, style and size; a table might be associated with the number of columns and rows, *etc.* Phase 2 consists of *block classification*. The result of phase 2 is an assignment of labels (title, regular text, picture, table, *etc.*) to all the blocks using properties of individual blocks from phase 1, as well as spatial layout rules. Phase 3 consists of *logical grouping* and ordering of blocks. For OCR it is necessary to order text blocks. Also the document blocks are grouped into items that "mean" something to the human reader (author, abstract, date, *etc.*), and is more than just the physical decomposition of the document. The output of phase 3 is a hierarchical tree-of-frames, where the structure is defined by the shape of the tree and the content is stored entirely in the leaves. The tree-of-frames can be converted to the SGML representation to ensure portability, easy information retrieval, editability and efficient semantic indexing of the document.

3.1 Block segmentation

Approaches for segmenting document image components can be either top-down or bottom-up. Top-down techniques divide the document into major regions which are further divided into sub-regions based upon knowledge of the layout structure of the document. Bottom-up methods progressively refine the data by layered grouping operations.

3.1.1 Top-down methods

Smearing is based on the Run-Length Smoothing Algorithm (RLSA) [5]. It merges any two black pixels which are less than a threshold apart, into a continuous stream of black pixels. The method is first applied row-by-row and then column-by-column, yielding two distinct bit maps. The two results are then combined by applying a logical AND to each pixel location. The output has a smear wherever printed material (text, pictures) appears on the original image.

The *X-Y tree* method assumes that a document can be represented in the form of nested rectangular blocks [6]. A "local" peak detector is applied to the horizontal and vertical "projection profiles" to detect local peaks (corresponding to thick black or white gaps) at which the cuts are placed; it is local in that the width is determined by the nesting level of recursion, *e.g.*, gaps between paragraphs are thicker than those between lines.

The *Hough transform* approach exploits the fact that documents have significant linearity. There exist straight lines in tables and diagrams. Centroids of connected components corresponding to text also line up. Columns of text are separated by straight rivers of white space. Text itself can be viewed as thick textured lines. The Hough transform is a technique for detecting parametrically representable forms like straight lines in noisy binary images. In fact, we have shown that the Hough transform is truly a representation of the projection profiles of the document in every possible orientation [7]. The accumulator array can be checked for the particular orientation which has the maximum number of transitions to and from a white pixel value. Transitions corresponding to text are usually regular and uniform in width and thus are easy to identify. Maximum values are registered at the center line of the characters and slightly lower values corresponding to the ascender and descender lines. The analysis of the accumulator array has an added advantage. It can provide the angle of skew in the document. Since the segmentation phase assumes that blocks align parallel to the page margins, skew detection and correction is an essential preprocessing step.

Activity identification is a method developed at CEDAR for the task of locating destination address blocks on mailpieces. The goal is to use a simple technique for documents with simple layout structure (*i.e.*, small number of blocks). It quickly converges on an area of maximum activity by identifying regions with a high density of connected components. The method is useful for goal-oriented tasks like finding the destination address block on a facsimile cover page which is predominantly empty.

3.1.2 Bottom-up methods

In bottom-up segmentation methods, the image is first processed to determine the individual connected components. At the lowest level of analysis, there would be individual characters and large figures. In the case of text, the characters are then merged into words, words are merged into lines, lines into paragraphs *etc.*

The "docstrum" method of image segmentation [8] uses K-nearest-neighbor clustering of connected components. For each connected component in the image, K nearest neighbors are found by investigating the neighborhood around it. A neighbor is defined by the distance between the centroids of the component and the neighbor, as well as the angle to the horizontal of the line connecting the centroids. Then, by analyzing these distances and angles for the whole image, values of the skew, between-text-line spacing, and within-text-line spacing are obtained. These parameters are then used to merge the connected components into text lines, and text lines into text blocks. The main advantages docstrum has over other bottom-up approaches to image segmentation are its relative independence from *a priori* knowledge of layout rules and structure, and the concurrent determination of skew and block layout. This independence of *a priori* knowledge makes it possible for *docstrum* to operate over a wide range of documents.

The application of different operators for bottom-up grouping can be coordinated by a rule-based system. Most characters can be quickly identified by their component size. However, if characters touch a line, as is often the case in tables and annotated line-drawings, the characters have to be segmented from the lines. One technique for segmenting characters from line structures is to determine the high neighborhood line density (NLD) areas in the line structures. Figure 1 shows an example of the type of rules that can be used.

3.2 Block Classification

Blocks determined by the segmentation process need to be classified into one of a small set of predetermined document categories. Knowledge of the layout structure of a document can aid the classification process. For instance, if it is known *a priori* that a given document is a facsimile cover page, then inferences like "the central block must be labeled as the destination address and the top of the document must be labeled as the name of the organization", *etc.*, are plausible.

A statistical classification approach using textural features and feature space decision techniques has been developed at CEDAR to classify blocks in newspaper images [9]. Two matrices, whose elements are frequency counts of black-white pair run lengths and black-white-black combination run lengths, are used to derive the texture information. Three features are extracted from the matrices to determine a feature space in which labeling of blocks as picture, regular text, headlines (major and minor), and annotations (*e.g.*, captions to photographs) is accomplished using linear discriminant functions.

3.2.1 Text determination

A method of associating labels like line-of-characters, paragraph, title, author, text columns, picture, table, line-drawing, *etc.*, to the blocks can be done on the basis of their spatial structure and relative and absolute positions in the document. The following are rules that can be implemented in a knowledge-based system:

line-of-characters:	sequence of adjacent character blocks of pre-determined height
paragraph:	sequence of blocks of line-of-characters of same length
column:	large blocks with roughly equal frequency of black and white pixels

3.2.2 Discriminating between handwritten and machine-printed text

Separating machine-printed text from handwritten annotations is necessary for invoking appropriate recognition algorithms. One method that performs this discrimination well with postal addresses is based on computing the histogram of heights of the connected components. Handwritten components tend to have a wider distribution in heights than print.

3.2.3 Figure classification

A figure block can belong to one of the following categories: half-tone picture, line-drawing (diagrams), and table. Pictures are gray-scale images and can be separated from tables and diagrams (binary images) by a histogram analysis of the gray-level distributions. Although both diagrams and tables predominantly comprise of straight lines and interspersed text, the fact that the straight lines in tables run only vertically and horizontally can be used to advantage. Analysis of the Hough transform accumulator array is used to separate tables from diagrams.

3.2.4 Goal-oriented document segmentation and classification

In certain document segmentation tasks it is only necessary to extract a given block(s) of interest. An example of this is locating an address block on a mail piece [10]. Several top-down and bottom-up tools can be used to segment candidate blocks. They are resolved by a control structure to determine the destination address block by considering spatial relationships between different types of block segments.

3.3 Logical Grouping

It is necessary to provide a logical ordering/grouping of blocks to process them for recognition and understanding. Textual blocks corresponding to different columns have to be ordered for performing OCR.

Journal pages can have complicated layout structures. They are usually multi-columned and can have several *sidebars*. Sidebars are explicit boxes enclosing text and figures used for topics that are not part of the mainstream of text. They can either span all or some of the columns of text. Readers wanting a quick overview usually read the main text and avoid the sidebars. The block classification phase labels both the mainstream of text and the sidebars as textual blocks. It is up to the logical grouping phase to order the blocks of the main text body into a continuous stream by ignoring the sidebars.

Another grouping task pertinent to this phase is matching titles and "highlight" boxes to the corresponding text blocks in the mainstream. When a title pertains to a single-columned block, associating the title with the corresponding text is straightforward. However, titles can span several columns, and sometimes be located at the center of the page without aligning with any of the columns, making the task of logical grouping challenging.

One method of logical grouping using rules of the layout structure of the document has been developed at CEDAR for newspaper images ([11], [12]). Figure 2 shows some examples of rules derived from the literature on "Newspaper design" used in the spatial knowledge module of this approach.

Blocks with different labels (photograph and text) that are not necessarily adjacent might have to be grouped together. For instance, a photograph and its accompanying caption together form a logical unit and must be linked together in the output representation.

3.4 Output representation

The layout structure of a document divides and subdivides the document into physical rectangular units, whereas the logical structure divides and subdivides the document into units that "mean" something to the reader (*e.g.*, a newspaper story or a journal article). The output representation of the document decomposition module is a tree-of-frames where each node in the tree is a logical block represented by a *frame* and the root of the tree corresponds to the entire document. The shape of the tree reflects the structure of the document and the contents of the document is stored entirely at the leaves of the tree. A frame is an *attribute, value* representation that consists of slots where each attribute can store a value. For instance, a block of text can have the attributes of font type, size, style and a flag to indicate if it is handwritten or machine printed. The OCR benefits by knowing the values of these attributes prior to recognition.

Standard Generalized Markup Language (SGML) provides the syntax (in terms of regular expressions) for describing the logical structure of documents. The tree-of-frames representation can be transformed into SGML syntax without any loss of information. SGML parsers are available that check the validity of the syntax of the logical structure and evaluate the correctness of the decomposition of the document. SGML also provides *link* mechanisms that can reproduce the entire document from the syntax rules. Such a bidirectional mapping between a document and its representation provides a means for semantic indexing and efficient information retrieval.

4 A system for document structure analysis

At CEDAR we have developed and implemented a knowledge-based system for document decomposition and structural analysis that uses a rule-based model. This system, called DeLoS ("DErivation of LOgical Structure") takes as input the digitized image of a newspaper page and produces as output a symbolic description of the logical structure of the page, including a classification of each of the printed blocks on the page, a grouping of these blocks into logical units, and the reading order of text blocks within each of these units. Various image processing operations are performed to extract basic properties of the components of the image, and this data is analyzed under the control of a rule-based system, which operates in conjunction with a global data structure to monitor the entire classification, grouping, and block-ordering process. Some details of the design and development of this system have been described in previously published papers ([13], [14]).

4.1 System Organization

Figure 3 shows the components of the DeLoS system. The unbroken arrows between the boxes indicate the flow of data between the various components of the system, and the dotted arrows represent the control flow in the system.

As seen in Figure 3, the DeLoS system consists of a multi-level, rule-based reasoning system, an image processing sub-system, and a partitioned global data structure. The rule-based system utilizes a top-down, backward-chaining structure. An inference engine within the rule-based system makes deductions about the document using a hierarchical knowledge base that contains rules describing all the identifiable characteristics of document images. The rules are classified into three levels: knowledge rules, control rules, and strategy rules. The global data structure facilitates the transfer of information between the image processing modules and the rule-based system. A common data area stores all intermediate computation results and other control information.

The image-processing modules directly access the document image to extract various kinds of information about the document. Intrinsic properties of the different printed blocks as well as the spatial relationships between the different blocks constitute the information that is passed back to the control structure through the global data structure. Some image-processing modules perform the basic tasks of binarization, connected-component analysis, etc., so as to prepare the image for other modules (e.g., block segmentation) that can extract information from it.

4.2 Rule-Based Control Structure

The rule-based system consists of three levels of rules. The three-level rule structure for this system was inspired by research on low-level image segmentation of natural scenes using a rule-based system [15]. It was demonstrated there that using a hierarchical structure of three progressively abstract levels of rules provided a large amount of flexibility in the inference mechanism, and allowed a modular formulation of the solution within the image analysis problem domain. The three levels into which the rules in this system are classified are as follows:

- Knowledge Rules (level 1)
- Control Rules (level 2)
- Strategy Rules (level 3)

The domain knowledge in the system is maintained in the Knowledge rules, and the control decisions in the system are made by the Control and Strategy rules, as described below.

4.2.1 The Knowledge Base

The Knowledge rules comprise the knowledge base that contains all the domain knowledge for the system, expressed in terms of first-order predicates. These rules define the general characteristics expected of the usual components of a document image and the usual relationships between such components in the image.

Thus, all common characteristics of different types of document blocks (e.g., text blocks, photographs, etc.), as well as spatial constraints commonly followed in document layout (e.g., the positioning of captions relative to photographs, etc.), are encoded into the knowledge base. These knowledge rules in the knowledge base contain all the properties and spatial relations for different types of document blocks, and can be used for block classification, block grouping, or text block ordering as and when required according to the control strategy.

Knowledge rules can be further categorized into the following:

- *Unary rules:* These rules allow for the classification of a block based solely upon its properties; *e.g.*, a thin graphical block in a newspaper image can be classified as a row separator or a column separator based solely upon its size and aspect ratio.
- *Simple binary rules:* These rules depend on the positions of adjacent blocks for classification; *e.g.*, a half-tone block in a newspaper image can be classified as a page banner based upon not only its size and aspect ratio, but also its position relative to its neighboring blocks.
- *Complex binary rules:* These rules depend on the identities of adjacent blocks for classification; *e.g.*, a small-text block in a newspaper image can be classified as a caption based not only on its own size and aspect ratio, but also the properties and identities of its neighboring blocks (i.e., whether it is adjacent to a photograph block and meets certain spatial constraints).

The hierarchical structure of the knowledge base mentioned above is thus created by these different categories of rules. Some examples of knowledge rules are shown in Figure 4.

4.2.2 The Inference Engine

The control structure for the rule-based system contains an inference engine which is also rule-based, and contains two levels of rules: *control rules* and *strategy rules*. These rules regulate the analysis of the document image, and decide when a consistent interpretation of the image has been obtained.

The rules comprising the inference engine are also formulated in terms of first-order predicates. The control structure determines the order in which these rules are executed in order to test various conditions effectively. Control rules regulate the invocation of the knowledge rules, based on appropriate data configurations or processing states. Control rules can be further categorized as follows:

- *Focus-of-attention Rules:* Focus-of-attention rules are designed to recognize data configurations that indicate the blocks on which to apply knowledge rules next, e.g., a rule that decides that the identity of the block directly above the current block should determined next.
- *Meta-Rules:* Meta-rules, on the other hand, are designed to recognize processing states and data configurations that indicate what category of knowledge rules are to be executed next as the highest priority, e.g., a rule that decides that all the relevant knowledge rules for captions should be executed next so as to determine whether the current block is a caption.

Strategy rules guide the search in a more general way, i.e., they determine what control strategy is to be followed at any given time for analyzing the image. This means that the strategy rules regulate the invocation of, and determine the execution order of, the control rules. Strategy rules also decide on the stopping criteria for the system, i.e., whether a consistent interpretation and grouping of the blocks in the document image has been achieved (as determined by the absence of incomplete block or unit data in the global data structure, and by the completeness of the logical structure tree). Therefore, there is a set of strategy rules for block classification, another set for block grouping, and yet another for text block ordering.

Examples of control rules are shown in Figure 5. Examples of strategy rules are shown in Figure 6.

4.3 The Global Data Structure

The global data structure stores the physical structure and logical structure information for the document being processed. It also facilitates the transfer of information between the rule-based system and the image processing modules. A common data area stores all intermediate computation results and other control information, and provides the framework for the construction of the trees representing the document structures.

In this system, the global data structure is divided into the following partitions:

- *Domain Data Partition:* This is the major part of the data structure, and contains the data and inferences about the image blocks and units. Thus, this is the partition in which the frame-trees representing the physical and logical structures are formed and modified.

- *Control Data Partition:* This contains information about the status of the classification, grouping and read-ordering process as well as details about any partial data or results being maintained in the global data structure.

The domain data partition contains the domain data (i.e., data and inferences about the image blocks and units) that is transferred by the image processing routines to the rule-based inference engine. It also contains the physical and logical structures, which are organized as a set of block frames and unit frames in a hierarchical form.

The control data partition in this system is logically separate from the domain data partition, and contains flags, switches, and other variables which indicate exactly what data currently exists in the domain data partition, what part of the document image was processed to obtain that data, what operations were performed, and with what parameters. It also contains information about the status of the overall classification and grouping process, and also details about any partial data or results that are being maintained in the global data structure.

4.3.1 Representation of Domain Data

As mentioned above, the domain data partition contains data that has been extracted from the document image by the various image-processing routines. A common representation for this data is essential for the purpose of access and modification by the different components of the system. The representation chosen must reflect the original structure of the document and at the same time be flexible enough for constant modification, since data is obtained at various times during the block classification, grouping, and ordering process.

The input image data is represented initially as a list of blocks with their physical properties and their type. The physical structure tree is created from this list of blocks. Each block in the image data is represented by a frame. A frame is an attribute-value data structure that consists of a number of attribute "slots", to each of which one can associate a value. Thus, each attribute slot in a frame represents a particular characteristic of a block or unit. In our system, no slot can contain a default value; each value maintained in a frame must be explicitly placed there.

There are two kinds of frames:

1. *Block frames:* These frames represent physical characteristics of document blocks and are thus structural in nature.

2. *Unit frames:* These frames are more conceptual in nature, and represent the logical groupings of the different document blocks.

Each block frame contains data for one block in the image. Different characteristics of a block are represented as slots in a block frame. The different slots in the current DeLoS system are shown in Figure 7 (a).

Unit frames represent logical groupings of the blocks in the image (e.g., a newspaper story, consisting of a set of text paragraphs, a photograph and its caption, and a headline). Information in a unit frame can only be filled in after the classification of all the blocks in the document image has been made and the blocks have been grouped into logical units. The slots in a unit frame are shown in Figure 7 (b).

Unit frames and block frames, when the slots are all filled in, make up a tree structure of frames, since each unit is a parent to several blocks or other sub-units. It is possible to put a unit number into the "list of block numbers" slot of a unit frame, which makes that unit a child unit of the current one. Thus, a tree structure is created, whose root node is the unit frame representing the entire document page, and all other nodes are either unit frames representing logical subdivisions of the page, or block frames representing physical blocks of printed matter in the page. The above representation therefore allows us to specify logical sub-units of a given unit. For example, a given story on a particular topic with one major headline may have sub-stories under separate minor headlines. The system, by allowing unit frames to be child nodes of other unit frames, allows us to efficiently represent such a hierarchy among the units in a document page.

4.4 System Implementation

The DeLoS system has been implemented on a Sun SPARC-2 workstation running SunOS. Document images are scanned using a flatbed scanner at a resolution of 300 pixels per inch, and the resulting bitmap of the document in Sun raster format is converted into HIPS format [16] prior to any further processing.

Because of the propensity of different languages to allow efficient representation and manipulation of certain types of information, sub-systems in DeLoS are implemented in different programming languages. The image processing routines, which provide the inputs to the system and which require efficient and fast interaction with the actual image data, are implemented in C. The rule-based control structure, the knowledge base, and the global data structure are implemented in Prolog. The interaction between these sub-systems is handled via system routines (written in C and Prolog) which perform tasks such as translation of formats, consolidation of image data, and document-structure reconstruction.

The acceptable codes for block types (block categories) are given in Figure 8 (a). The acceptable codes for block labels (block classifications) for newspapers are given in Figure 8 (b). As can be seen from the list, the first digit of each block label refers to its block type, thus making it easier for the system to keep track of which block types can be legitimately translated into which block labels.

4.5 System Inputs

Document pages are scanned in as 300 ppi gray-scale images. After conversion to HIPS format, each image is then binarized, filtered to remove noise, and corrected for skew. Binarization is done using an adaptive thresholding method. Next, connected component analysis is done on the binary image. Segmentation (using "docstrum") is then performed, and text/graphics discrimination is performed. From the block-segmentation information, a block pre-processing program performs block-type categorization. The output of the block-type categorization program consists of a list of the blocks in the document showing their coordinates and type. This block list is then input into the logical structure derivation module.

4.6 System Outputs

The primary output from the DeLoS system is a data representation of the logical structure of the document, i.e., a tree showing all the classified blocks in the document image, giving all the relevant extracted feature details for each block, and an indication of the logical unit to which the particular block belongs, as well as an ordered listing of all the text blocks in each unit. This data structure is output in the form of an editable text file that contains information about the identity, position, and other relevant properties of all the logical blocks in the image, as well as the logical groupings of these blocks into logical units. In addition to this, the output also contains pointers between blocks to represent the reading order of the text blocks within each unit.

The data structure showing the different classified blocks and their logical groupings in a document is eminently suitable for transmission to another location, where the document can be reconstructed from its description. Also, such a data structure can be used by a higher-level interactive system that accepts user queries and attempts a semantic interpretation of the document, including text reading and natural language understanding.

4.7 Image Processing

The image processing modules directly access the document image to extract various types of information about the document, primarily the intrinsic properties of the printed blocks and the block type for each block. This information is stored in the global data structure and used as input data for the rule-based system.

The basic image processing operations performed on the scanned gray-level document image include binarization, connected component analysis, image segmentation, and block categorization. Figure 9 shows the steps in the transformation of the original document into the format required for input into the DeLoS system.

The original, gray-level document image obtained by scanning the physical document is first converted from the original Sun raster format to the HIPS format. It is then binarized using an adaptive thresholding algorithm [17]. Then, connected component analysis [18] is performed on the binary image. Next, segmentation is performed on the image from the connected component data, by using the "docstrum" method.

As previously mentioned, the "docstrum" method of image segmentation uses K-nearest-neighbor clustering of connected components. Because the "docstrum" program relies on neighborhood clustering, it has to be run several times, each time removing the identified text blocks from the connected components list. This prevents the large text blocks (e.g., major headlines) from being type-categorized as graphics, which is what happens if "docstrum" is only run once on the entire image. In order to allow a more accurate analysis of the larger components, the connected component size thresholds in "docstrum" are changed appropriately after the removal of the smaller connected components at the end of each run.

The result of segmentation is that each of the printed blocks on the image is isolated and categorized as text or graphics. Further processing of the blocks information, using block size and connected component characteristics, results in the categorization of each block into a *block type*, i.e., "small-text", "medium-text", "large-text", "line-drawing", or "half-tone". The list of all the blocks in the image, giving the coordinates of their enclosing rectangle as well as the basic block type, is the input to the knowledge-based document logical structure derivation module.

An alternative approach to block-type categorization that was originally used by the DeLoS system is the method of texture analysis [9], because text areas can be considered to have some common textural characteristics. For this, homogeneous rectangular blocks are first segmented out of the image using methods such as run-length smoothing and recursive horizontal/vertical cuts. The type categorization approach is based on statistical textural features and feature space decision techniques. Two matrices, whose elements are frequency counts of black-white pixel pair run-lengths and black-white-black combination run-lengths, are used to derive texture information. Three features are extracted from the matrices to determine a feature space in which block type categorization is accomplished using linear discriminant functions. The feature spaces created by these statistical features are effective in categorizing blocks into a set of predetermined document block categories.

5 Experimental Results

The DeLoS system has been trained using data from page images of *The Buffalo News*, and extensively tested on data from digitized images of pages from *The Buffalo News* and *USA Today* newspapers.

Figure 10 and Figure 11 show the results obtained from the image of a page from *USA Today.* Figure 10 (a) shows the original newspaper image. Figure 10 (b) shows the segmented image, with the segmented blocks shown in white. Figure 11 shows the output of the DeLoS system. Each segmented and type-categorized block is represented by two lines. The details shown for each block are the block number, the block dimensions (width, height), the block coordinates (x,y coordinates of the upper-left corner), the block type, the block label, the previous and next text block in logical reading order, and the unit number to which the block belongs. This output contains the filled-in information about the blocks (i.e., block labels) and also unit details (unit number, number of blocks, and a list of blocks in the unit). A unit number of 99 represents a block that was not grouped into any specific unit. The filled in block numbers in the "prevtxt" and "nexttxt" fields create the logical links that determine the read-order of text blocks within a unit. A "prevtxt" value of 99 means that the current block is the first text block in its unit in terms of read-order, and a "nexttxt" value of 99 means that the current block is the last text block in its unit in terms of read-order. For non-text blocks, the "prevtxt" and "nexttxt" fields have a null value (i.e., block number "0").

Figure 12 shows the reconstructed physical structure of the page, showing the different blocks in the image. The reconstructed physical structure is obtained by running a format conversion program to convert the DeLoS output to a LaTeX format file (consisting primarily of commands in the LaTeX *picture* environment), which is then processed by LaTeX to produce the output structure shown.

Evaluation of the results on each image was done by comparison with a truth file that was generated from the original image. This truth file contained the true classification, grouping and read-ordering data.

So far, the DeLoS system has been tested on a variety of pages from *The Buffalo News* and *USA Today*. A total of 44 images of newspaper pages (12 images of pages of *USA Today* and 32 images of pages of *The Buffalo News*) were analyzed through the system.

Overall, the DeLoS system performed fairly well for images from *The Buffalo News* as well as *USA Today*. Figure 13 shows the performance of the system for the *USA Today* newspaper pages, in terms of percentages of the original blocks correctly classified, grouped and read-ordered. It also shows the percentages of correctly segmented blocks that are correctly classified, grouped and read-ordered. Performance results for *The Buffalo News* followed a similar pattern (and have been described in detail in [11]).

As we can see from the table in Figure 13, block segmentation and block type-categorization in the original images proved to be a deciding factor in the performance of the system. The "docstrum" program when run on the original binary images gave only the text blocks accurately, and categorized most of the headline blocks as graphics. The connected components comprising these blocks were then extracted and "docstrum" was re-run on only these components, which then yielded a more accurate categorization of medium-sized text blocks (corresponding to minor headlines). Another similar extraction of connected components and a re-run of "docstrum" yielded the large-sized text blocks (corresponding to major headlines). Also, since graphical blocks often have similar geometrical extents as rectangular enclosing rectangles, the latter were categorized by computing the block's pixel-density $pixden$, which is defined as the ratio of the number of pixels in the block to the area of the block. A very low value of $pixden$ indicates an enclosing rectangle. In addition, horizontal and vertical lines were categorized by recognizing blocks which had a very small height or width respectively.

6 Summary

Document decomposition & structural analysis, one of the major processes needed to develop a DIU system, has been described. Decomposition and structural analysis techniques are responsible for decomposing a document into several regions, each of which contains homogeneous entities. These regions are then grouped into logical units to form a high-level interpretation of the document structure. A system that performs document decomposition and structural analysis (including logical grouping and read-order determination) has been presented, along with results obtained by applying the system to images of newspaper pages.

We have shown through the development of the model and the results obtained from the implementation of the DeLoS system that:

- The geometric structure of a document page can be translated into rules.
- The logical structure of different types of documents can similarly be encoded into rules.
- The logical structure of a document page can be derived from the physical structure by the application of these structural and semantic rules.

The most obvious use of the DeLoS system is as a part of a comprehensive document understanding system that uses the logical structure information output by this system to read the information in the individual units in a document. Applications include document browsers and document library archival systems.

Digital document libraries have recently gained tremendous national attention, mainly because of the U.S. Government's efforts at creating a *National Information Infrastructure*. The applications of DIU systems in creating intelligent digital libraries are manifold. Systems that can extract logically linked parts of a document and retrieve them on demand can significantly improve the usefulness and performance of document libraries.

6.1 Future Directions

The DeLoS system derives as complete a logical description of a document page as is possible without performing actual text reading and semantic understanding of the document contents. In fact, spatial information alone provides a significantly large amount of knowledge that enables an effective logical structure derivation. However, a more complete understanding of document structure and content can obviously be obtained by involving text understanding in the logical structure derivation process. Therefore, any further work based on the current document logical structure derivation model needs to include text understanding.

Other potential extensions to the DeLoS system that are suitable to be pursued for future research work include *multi-page logical grouping, linking non-adjacent blocks to logically grouped units,* and *semantic analysis for read-order determination.*

References

[1] S.N. Srihari. Document Image Understanding. *IEEE Fall Joint Computer Conference*, Dallas, Texas, 1986, pp 87-95.

[2] G. Nagy. What Does a Machine Need to Know to Read a Document? *Proceedings of the Symposium on Document Analysis and Information Retrieval*, Las Vegas, NV, 1992, pp 1-10.

[3] S.W. Lam, A.C. Girardin and S.N. Srihari. Gray-Scale Character Recognition Using Boundary Features. *SPIE/IS&TSymposium on Electronic Imaging Science &Technology*, San Jose, California, 1992.

[4] Y.Y. Tang, C.Y. Suen, C.D. Yan and M. Cheriet. Document Analysis and Understanding: A Brief Survey. *First International Conference on Document Analysis and Recognition*, Saint Malo, France, 1991, pp 17-31.

[5] K.Y. Wong, R.G. Casey and F.M. Wahl. Document Analysis System. *IBM J. Res. Develop.* **26**, No. 6, 1982, pp 647-656.

[6] G. Nagy, S.C. Seth and S.D. Stoddard. Document Analysis with an Expert System. In *Proceedings of Pattern Recognition in Practice II*, Amsterdam, June, 1985.

[7] S.N. Srihari and V. Govindaraju. Textual Image Analysis Using the Hough Transform. *International Journal of Machine Vision and Applications*, **2(3)**, 1989, pp 141-153.

[8] L. O'Gorman. The document spectrum for page layout analysis. In *IAPR International Workshop on Structural and Syntactic Pattern Recognition*, Bern, Switzerland, August 1992.

[9] D. Wang and S.N. Srihari. Classification of newspaper image blocks using texture analysis. *Computer Vision, Graphics, and Image Processing*, **47**, 1989, pp 327-352.

[10] S.N. Srihari. Feature extraction for locating address blocks on mail pieces. *From Pixels to Features*, J.C. Simon (ed.), Elsevier Science Publisher B.V. North-Holland, 1989, pp 261-273.

[11] D. Niyogi. *A Knowledge-Based Approach to Deriving Logical Structure from Document Images.* PhD thesis, State University of New York at Buffalo, 1994.

[12] V. Govindaraju, S.W. Lam, D. Niyogi, D. Sher, R.K. Srihari, S.N. Srihari and D. Wang. Newspaper Image Understanding. *Lecture Notes in Artificial Intelligence*, Vol. 444, J. Siekmann (editor), Springer Verlag, NY, NY. 1989, pp 375-386,.

[13] D. Niyogi and S.N. Srihari. A rule-based system for document understanding. In *Proceedings of AAAI-86*, vol. 2, pp 789–793, Philadelphia, PA, August 15–22, 1986.

[14] D. Niyogi and S.N. Srihari. Knowledge-based derivation of document logical structure. In *Fifth International Conference on Document Analysis and Recognition (ICDAR '95)*, pp 472–475, Montreal, Canada, August 14-16, 1995.

[15] A.M. Nazif and M.D. Levine. Low level image segmentation: An expert system. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(5):555–577, September 1984.

[16] M. Landy, Y. Cohen, and G. Sperling. HIPS: A UNIX-based image processing system. *Computer Vision, Graphics, and Image Processing*, 25:331–347, 1984.

[17] Y. Liu, R. Fenrich, and S.N. Srihari. An object attribute thresholding alogrithm for document image binarization. In *Proc. of ICDAR-93*, pp 278–281, Tsukuba Science City, Japan, Oct. 20–22, 1993.

[18] G.L. DeSilva. Syntax-directed text segmentation in page layout analysis. Master's thesis, State University of New York at Buffalo, 1993.

if connected component size is larger than a threshold
then the connected component is a figure (with likelihood L_i)

if neighborhood line density(NLD) of a figure is high
then the high NLD area is a character area (with likelihood L_j)

Figure 1: Example of rule using neighborhood line density.

Headlines occupying more than one printed line are left-justified.
Captions are below photographs, unless two or more photographs have a common caption.
An explicit box around a set of blocks signifies an independent unit.
Thin vertical lines are column separators.

Figure 2: Examples of rules used in newspaper structure analysis.

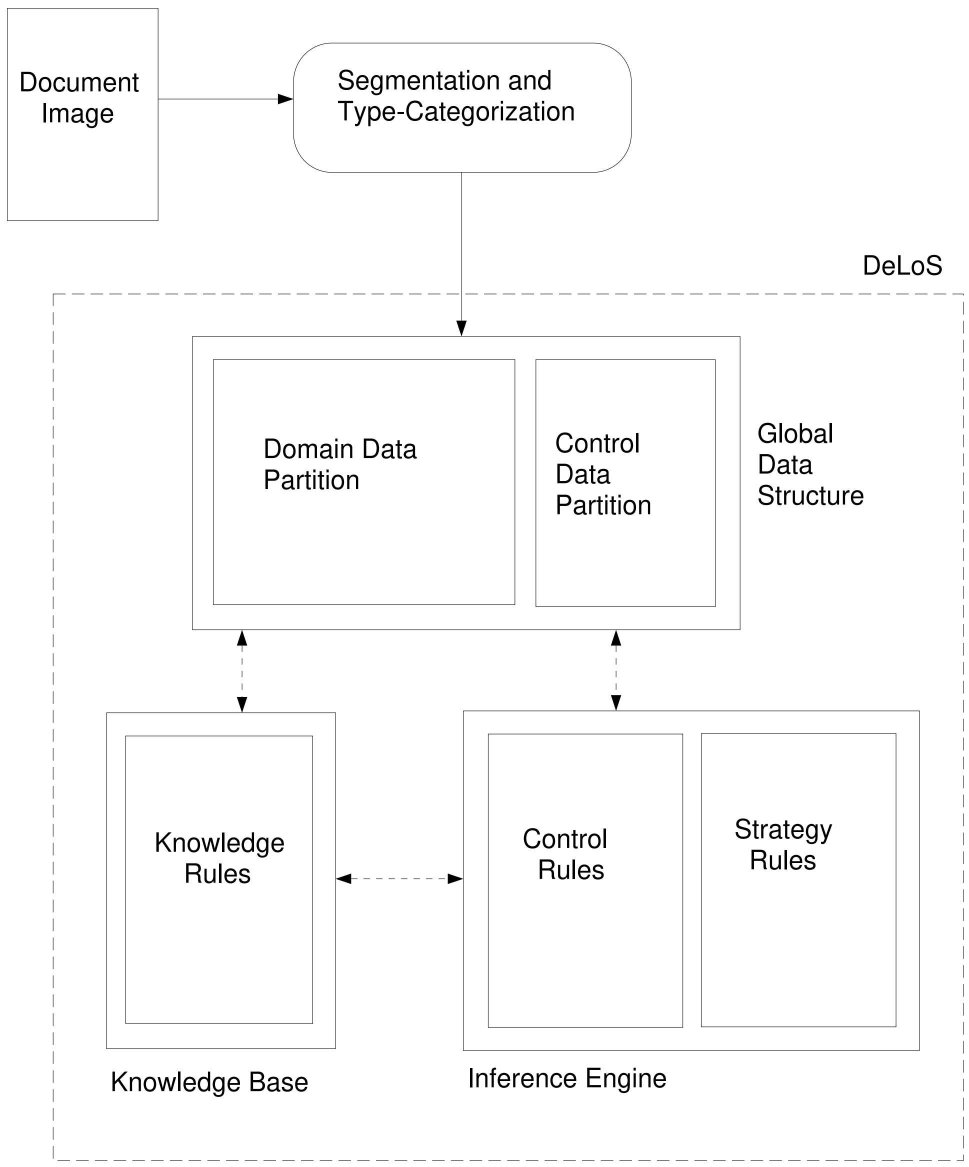

Figure 3: The DeLoS system.

```
IF a block X is of type ''small-text'',
  AND IF it is below another block Y which is a photograph,
  AND IF the widths of the two blocks are equal,
THEN block X is a photo-caption.


IF a block Z is of type ''large-text'',
  OR IF it satisfies the following three conditions:
      {it is of type ''medium-text'',
         AND it is below another block W,
         AND block W is not of type ''large-text'' or ''medium-text''},
THEN block Z is a major headline.
```

Figure 4: Examples of knowledge rules in the rule-based system.

```
[META-RULE]:

IF there are partial classification results for any blocks,
THEN select one (partially classified) block
  AND apply all classification rules for the selected block.


[FOCUS-OF-ATTENTION RULE]:

IF the grouping mode is on,
  AND a block has been selected,
THEN find all the immediate neighbors of the selected block.
```

Figure 5: Examples of control rules in the rule-based system.

```
IF classification of the selected block is not successful,
THEN store the partial results in the global data structure
  AND modify all appropriate indicators.


IF any partially grouped units remain,
THEN apply all unit-related control rules for each of these units
       until there are no more partial units.
```

Figure 6: Examples of strategy rules in the rule-based system.

```
(1) Block number
(2) Height
(3) Width
(4) Position (X--Y coordinates of top-left and bottom-right corners)
(5) Block Type (large-text/medium-text/small-text/photo/line drawing)
(6) Block Label
            [filled in by the system after classification]
(7) Unit to which it belongs
            [filled in by the system after grouping]
(8) Block numbers of the previous and next blocks in text reading order
            [filled in by the system after read-order determination]
```

(a) Block Frame

```
(1) Unit number
(2) Number of blocks contained in the unit
            [filled in and updated during grouping]
(3) Ordered list of block numbers
            [filled in during grouping, ordered during read-ordering]
```

(b) Unit Frame

Figure 7: Slots in block and unit frames.

```
1 - small text,     2 - medium text,      3 - large text,
4 - half-tone,      5 - line-drawing,     6 - other,
7 - vertical line,  8 - horizontal line,  9 - line rectangle.
```

(a) Block types

```
11 - text paragraph,    12 - caption,      13 - story byline,
21 - minor headline,
31 - major headline,
41 - photograph,        42 - page banner,  43 - thick horiz. band,
51 - line drawing,
61 - unknown,
71 - column separator,
81 - row separator,
91 - enclosing rectangle.
```

(b) Block labels for newspapers

Figure 8: Acceptable block types and block labels for the DeLoS system.

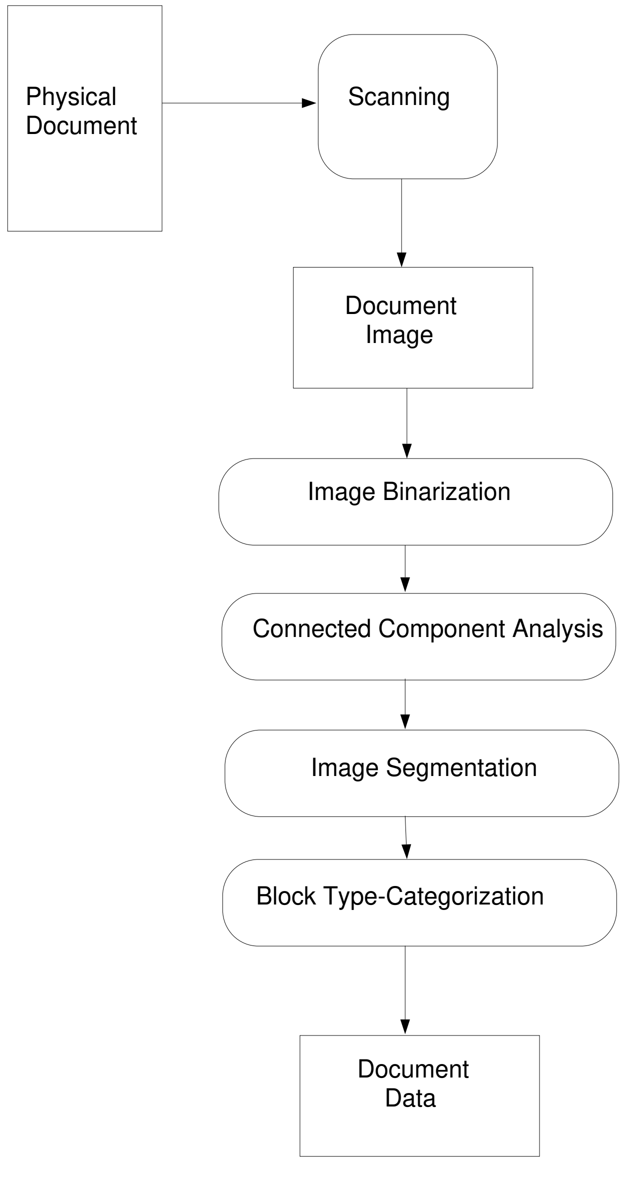

Figure 9: Image processing steps for documents.