

Performance Evaluation and Benchmarking of Six-Page Segmentation Algorithms

Faisal Shafait, Daniel Keysers, and Thomas M. Breuel

Abstract—Informative benchmarks are crucial for optimizing the page segmentation step of an OCR system, frequently the performance limiting step for overall OCR system performance. We show that current evaluation scores are insufficient for diagnosing specific errors in page segmentation and fail to identify some classes of serious segmentation errors altogether. This paper introduces a vectorial score that is sensitive to, and identifies, the most important classes of segmentation errors (over, under, and mis-segmentation) and what page components (lines, blocks, etc.) are affected. Unlike previous schemes, our evaluation method has a canonical representation of ground-truth data and guarantees pixel-accurate evaluation results for arbitrary region shapes. We present the results of evaluating widely used segmentation algorithms (x-y cut, smearing, whitespace analysis, constrained text-line finding, docstrum, and Voronoi) on the UW-III database and demonstrate that the new evaluation scheme permits the identification of several specific flaws in individual segmentation methods.

Index Terms—Document page segmentation, OCR, performance evaluation, performance metric.

1 INTRODUCTION

THE task of page segmentation is to divide the document image into homogeneous zones, each consisting of only one physical layout structure (text, graphics, pictures, ...). Therefore, the performance of optical character recognition (OCR) systems depends heavily on the page segmentation algorithm used. Over the last three decades, several page segmentation algorithms have been proposed in the literature (for a literature survey, please refer to that in [1], [2], [3]). In this paper, we present three main contributions to the state of the art in page segmentation:

1. Performance evaluation and comparison of six well-known algorithms for page segmentation using a state-of-the-art evaluation methodology [4]. We identify a severe flaw in the evaluation scheme when used for single column documents.
2. A novel, portable, and pixel-accurate representation for arbitrarily shaped page segments.
3. Several performance measures to identify and analyze different classes of segmentation errors made by a page segmentation algorithm.

The rest of the introduction section presents an overview of the state of the art in the above-mentioned areas and describes how our work augments the state of the art.

The problem of automatic evaluation of page segmentation algorithms is increasingly becoming an important issue

[5], [6]. Major problems arise due to the lack of a common data set, a wide diversity of objectives, a lack of meaningful quantitative evaluation, and inconsistencies in the use of document models. This makes the benchmarking of different page segmentation algorithms a difficult task. Recent page segmentation competitions [7], [8] address the need for comparative performance evaluation under realistic circumstances. However, a limitation of the competition-based approach is that competing methods only participate if they are implemented and used by a participant. It means several well-known algorithms might not be a part of the comparison at all.

The quantitative evaluation of page segmentation algorithms has received some attention in the past. An approach for measuring the quality of page segmentation algorithms by analyzing the errors in the text recognized by OCR was first proposed in [9]. However, text-based approaches have found little use since they measure the output of multiple steps and cannot be used to evaluate page segmentation alone. Yanikoglu and Vincent [10] presented a region-based page segmentation benchmarking environment, named Pink Panther. Their approach is based on representing regions as arbitrary polygons and, hence, becomes quite complex and cumbersome to use. Liang et al. [11] proposed a performance metric for document structure extraction algorithms by finding the correspondences between detected entities and ground truth. Das et al. [12] suggested an empirical measure of performance of a segmentation algorithm based on a graph-like model of the document. However, their performance measure does not support evaluation of non-Manhattan page layouts. Similar approaches have been presented for range image segmentation in [13] and for image segmentation in general [14]. Mao and Kanungo [4] presented an empirical benchmarking methodology based on text-line measure of page segmentation accuracy. This measure is particularly useful because it does not make assumptions about the layout of the

• F. Shafait and D. Keysers are with the Image Understanding and Pattern Recognition Research Group, German Research Center for Artificial Intelligence (DFKI GmbH), D-67663 Kaiserslautern, Germany.
E-mail: {faisal.shafait, daniel.keysers}@dfki.de.

• T.M. Breuel is with the Department of Computer Science, Technical University of Kaiserslautern, D-67663 Kaiserslautern, Germany.
E-mail: tmb@informatik.uni-kl.de.

Manuscript received 24 Apr. 2007; revised 10 Oct. 2007; accepted 5 Nov. 2007; published online 7 Dec. 2007.

Recommended for acceptance by L. O’Gorman and N. Lawrence.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log

Number TPAMI-2007-04-0246.

Digital Object Identifier no. 10.1109/TPAMI.2007.70837.

document. Besides, it requires only text-line-level ground truth. They have compared three research algorithms and two commercial products using this method. We extend the work by Mao and Kanungo [4] and add three more algorithms to the comparison (Section 2). The algorithms compared in [4] are x-y cut [15], docstrum [16], and the Voronoi diagram-based approach [17]. The algorithms added to the comparison in this work are the smearing algorithm [18], whitespace analysis [19], and the constrained text-line finding algorithm [20]. In addition, we also identify a limitation of the evaluation scheme when it is used for single-column documents and use a dummy segmentation algorithm to highlight the problem in such cases. An analysis of the errors coincides with our finding. We overcome this limitation by presenting a vectorial score that helps in judging the kind and extent of segmentation errors made by the analyzed algorithm (Section 2.3). This score is particularly useful in analyzing the behavior of a page segmentation algorithm for a given set of parameters. Our performance measures are based on a new way of representing layout information by embedding it in the color channels of a document image (Section 2.1). Such a representation allows convenient interchange of ground truth and segmentation results in terms of standard image formats.

We use the University of Washington III (UW-III) data set [21] for performance evaluation and benchmarking of the analyzed algorithms. The UW-III data set has a large number of documents with different degradation types and is one of the standard data sets for evaluating different document analysis tasks. The main strength of the data set is that, for each document, along with the ground-truth information for words, text-lines, zones, and ASCII text, a number of document and zone attributes are available. This makes the database suitable for quantitatively evaluating a wide variety of tasks related to document image analysis. Researchers have used the UW-III data set for evaluating their approaches related to different document analysis tasks like page segmentation [4], [11], block classification [22], layout analysis [23], table zone extraction [24], and document image classification [25]. We discuss the experiments performed and results obtained in Section 4, followed by the conclusion in Section 5. We have reported parts of the work presented in this paper in [26] and [27]. In this paper, we extend the performance measures presented in [27] by introducing the notion of correct segmentation. Additionally, we explain our methods in more detail than in [26] and [27], illustrating with visual examples where necessary. We have also included a correlation analysis of the errors made by each algorithm, which gives us interesting insights into the similarities of the behavior of different algorithms.

2 PERFORMANCE EVALUATION OF PAGE SEGMENTATION ALGORITHMS

The performance evaluation measure proposed in [4] is based on set theory. This measure is based on the assumption that a text block can be easily segmented into text-lines using horizontal projection. Let G be the set of all of the ground-truth text-line in a document image and $|G|$

denote the cardinality of the set G . Then, three subsets of text-lines are defined as follows:

1. The set of ground-truth text-lines that are missed (C), that is, they are not part of any detected text region.
2. The set of ground-truth text-lines whose bounding boxes are split (S), that is, the bounding box of a text-line does not lie completely within one detected segment.
3. The set of ground-truth text-lines that are horizontally merged (M), that is, two horizontally overlapping ground-truth lines are part of one detected segment.

The overall error rate is measured as the percentage of ground-truth text-lines that are not identified correctly:

$$\rho = \frac{|C \cup S \cup M|}{|G|}. \quad (1)$$

A ground-truth text-line is said to lie completely within one detected text segment if the area overlap between the two is significant. Significance is determined using two length thresholds in a number of pixels. The thresholds control the tolerance level along the horizontal and vertical directions such that differences in overlap less than the threshold in that particular direction are ignored.

Despite the many useful features, there is also a limitation of this approach. If a segmentation algorithm just takes the whole page as one segment, the split and missed errors vanish ($C = \emptyset, S = \emptyset$). Typically, for single-column documents, $M = \emptyset$. Hence, without doing anything, the segmentation accuracy can be high if there is a large proportion of single-column document images in the test data set. This effect was not considered in the original evaluation [4]. To check the severity of the problem, we have added a dummy segmentation algorithm into the comparison that returns the whole page as one segment, as discussed in Section 3.1.

We overcome this limitation by defining a vectorial score that clearly identifies the common classes of segmentation errors, including the undersegmentation problem identified above. This score is based on a new representation scheme for page segmentation described in Section 2.1. The vectorial score is described in Section 2.3.

2.1 Representation of Page Segments

Layouts of a document image are generally categorized into two main classes: *Manhattan* layouts and *non-Manhattan* layouts [1]. Manhattan layouts are defined as layouts that can be decomposed into individual segments by vertical and horizontal cuts. For Manhattan layouts, the individual zones can be represented by nonoverlapping rectangles. This representation is particularly useful due to its simplicity and segments of most of the structured documents, like technical journals or business letters, can be represented by their bounding rectangles. Therefore, this representation was adapted in the Document Attribute Format Specification (DAFS) format [28] used for representing the ground-truth zones for the UW-III data set. The DAFS format was developed with the intention of being used as a standard for the representation of document



Fig. 1. An example image to demonstrate color encoding of multiple layout levels. The top images show (a) word-level and (b) text-line-level segmentation representation, whereas the bottom images show (c) zone-level and (d) multiple layout-levels information encoded in different color channels of the same image.

images. However, it did not come into widespread use and other representations based on XML have emerged [29] for Manhattan layouts. For non-Manhattan layouts, the zones cannot be represented accurately by nonoverlapping rectangles. Instead, an XML-based representation of document zones by their bounding isothetic polygons was used in [7], [8]. A common problem with these approaches is that they need specialized software to view the files representing the page segmentation, thereby limiting their portability and ease of use.

To overcome these problems, we propose a new way of representing the page segments in color image format. Consider a document image decomposed into N homogeneous zones $Z_i, i = 1, \dots, N$. The document segmentation can be represented as an image in which each foreground pixel is assigned, as its value, the index of the segment Z_i to which it belongs. In practice, the pixel-based representation of page segmentation can be implemented as 24-bit RGB color images. This enables the use of up to $N = 2^{24}$ labels, which will be sufficient for virtually all images that are of interest. A particular color can be assigned to the page background (for example, 0xffffff) and to the noise pixels (for example, 0x000000). This representation of page segmentation is particularly convenient because it can be used to accurately represent different levels of layout in the same image, as shown in Fig. 1. Second, it is independent of the zone shape and it can be saved and exchanged using any lossless color image format.

2.2 Preparation of Pixel-Level Ground Truth

An image of a 300-dpi scanned A4 document usually contains over one million foreground pixels. The cost of coloring all foreground pixels using their respective segment label can be too high if all pixels are labeled individually. To overcome this problem, we consider two alternatives for preparing pixel-level ground truth.

1. A bounding polygon is drawn for each zone in the page image. The polygon is filled with a color representing the index of the zone contained inside the polygon (Fig. 2b). Then, each foreground pixel is assigned the color of the polygon that contains it (Fig. 2c). This approach is suitable when separation between zones in a page is significant. A benefit of preparing pixel-level ground truth with this approach is that a polygon of any shape can be drawn. For Manhattan layouts, a simple rectangle can do the task. For non-Manhattan layouts, a polygon can be quickly drawn around each zone. Hence, the cost of producing ground truth in this way is equal to the cost of producing any other bounding box-based ground truth in the case of Manhattan layouts. For non-Manhattan layouts, the cost for producing pixel-level ground truth can be much lower than other approaches because the polygons can be arbitrarily shaped and need not tightly enclose the containing zones.
2. If separation between page zones is not large, for instance, in the case of text-lines, the approach of

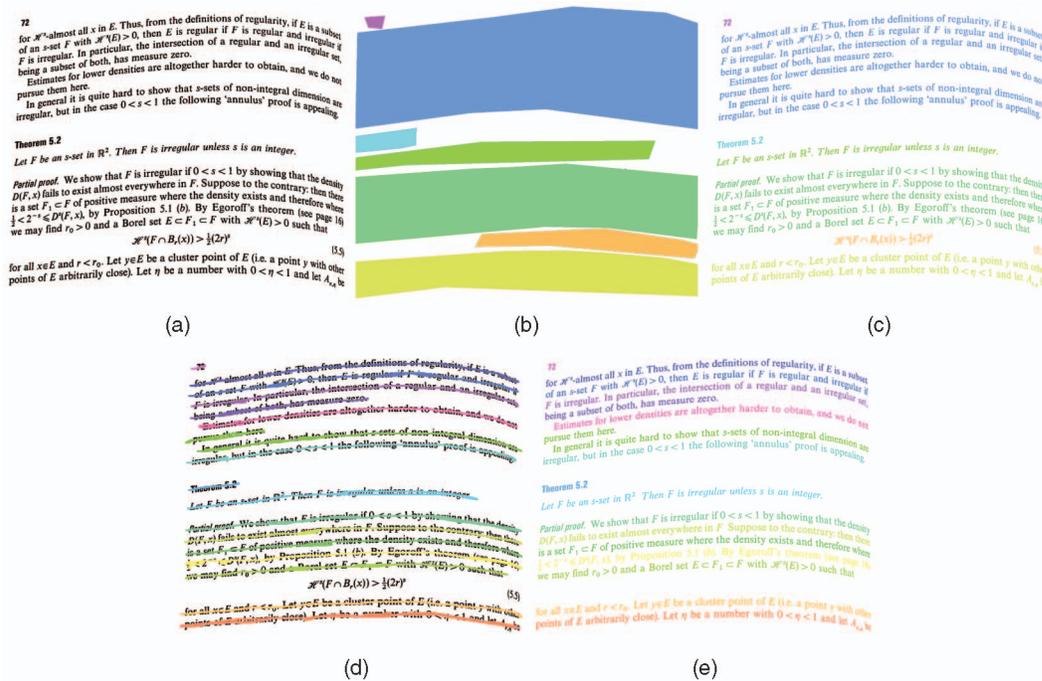


Fig. 2. An example image to demonstrate the process of generating pixel-level ground truth. The zone-level ground truth is prepared by first drawing a polygon around each zone (b) and then transferring the colors to the foreground pixels in the zone (c). The text-line-level ground truth is created by drawing lines (d) and then labeling the connected components touching these lines with the line color (e). (a) Original image. (b) Labeled zones. (c) Generated zone level ground truth. (d) Labeled text-lines. (e) Generated text-line ground truth.

creating ground truth with bounding polygons can become cumbersome. In such a situation, another approach can be taken. First, a line is drawn on a zone such that it touches or passes through all of the connected components of that zone (Fig. 2d). The color of the line is chosen to be the index of that zone. Then, connected components are extracted from the page and all of the foreground pixels in a connected component are assigned the color of the line that touches or passes through that component (Fig. 2e). In the final step, all small-sized components, like i-dots, punctuation marks, etc., are assigned the color of their closest neighbor if their distance to the closest neighbor is less than a threshold, chosen equal to x-height in our case. This step makes sure that any components that might not have been intersected in the first step get labeled as well.

Both of the above methods for creating pixel-level ground truth can be applied using any off-the-shelf image manipulation program like Gimp, MS-Paint, etc. These methods were applied in creating ground truth for the DFKI-1 warped documents data set used in the document image dewarping contest [30] held with CBDAR 2007.

2.3 Performance Evaluation

Based on the pixel-accurate representation of page segmentation, we define several performance measures to evaluate different aspects of the behavior of a page segmentation algorithm. Consider that we are given two segmentations in image form, the hypothesized segmentation H and the ground truth G . The images representing these segmentations should have the same dimensions and, for each corresponding pair of pixels in the two images, either both

pixels should belong to the background or both to the foreground. To compare the quality of a hypothesized segmentation against a ground-truth segmentation, we can construct a weighted bipartite graph called a *pixel-correspondence graph* [31] as follows: We associate with each color value in H or in G one node of the components in the graph, where the two components correspond to pixels of H and G , respectively. Since each segment has a unique color, each node represents a unique segment (either in H or in G). A segment that is labeled with a special color, like noise (see Section 2.1), can be removed at this stage. Then, an edge is constructed between two nodes such that the weight of the edge equals the number of foreground pixels in the intersection of the regions covered by the two segments represented by these nodes. If their corresponding segments do not overlap in H and G , no edge is needed.

If the hypothesized segmentation H agrees perfectly with the ground-truth segmentation G , then the pixel-correspondence graph will be a perfect matching. That is, each node in the two component of the graph has exactly one edge incident to it. If there are differences between the two segmentations, then the graph will not be a perfect matching. Instead, a node representing a segmentation in H or G may have multiple edges.

If P is the total number of pixels corresponding to one node (segment), M is the number of edges incident to that node, and $w_i, i = 1, 2, \dots, M$, is the weight associated with each edge, then $P = \sum_{i=1}^M w_i$. For each node on either component of the graph, w_i/P gives the fraction of pixels overlapping with each of its corresponding nodes.

An edge between two nodes is considered *significant* if $w_i/P \geq t_r$ or $w_i \geq t_a$, where t_r is a relative threshold and t_a is an absolute threshold. The use of t_r allows a tolerance in

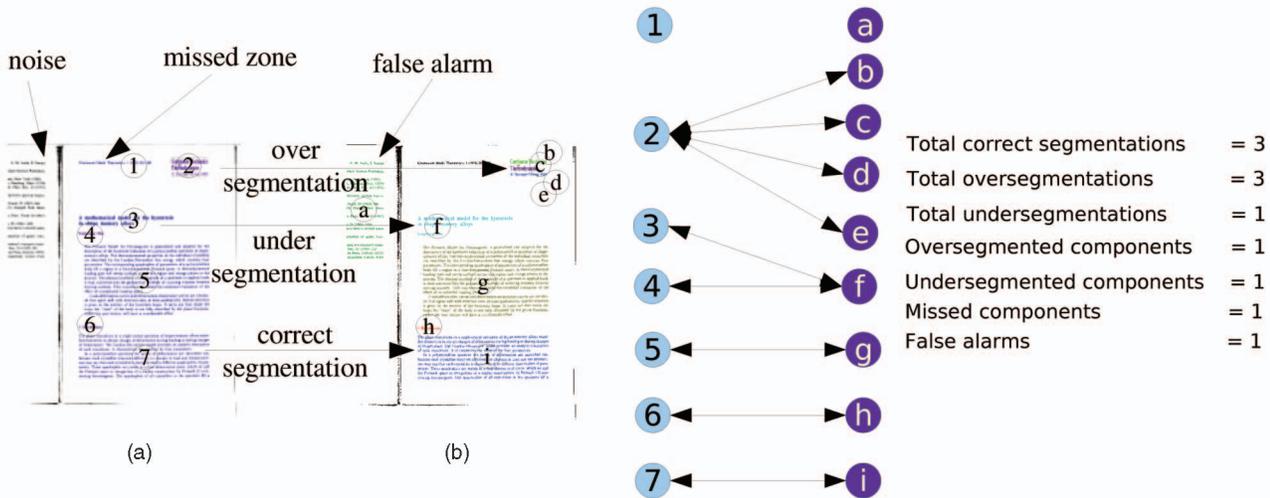


Fig. 3. Example image to illustrate different performance measures. The left image shows two color-coded document images. A pixel correspondence graph obtained from these images is shown on the right side. The nodes corresponding to the ground-truth segments are labeled 1-7, whereas the nodes in the segmented image are labeled a-i. Only significant edges are shown in the pixel correspondence graph. Based on the definitions given in Section 2.3, the values of each performance measure for this example are given on the right side of the graph. (a) Ground-truth image. (b) Segmented image.

the evaluation by ignoring fractional overlaps less than t_r . In practice, we have found $t_r = 0.1$ to be a good choice. However, if a segmentation algorithm completely fails and gives the whole page as one segment, regions containing less than 10 percent of the foreground pixels may get ignored. Therefore, an absolute threshold t_a is used to ensure that overlaps of more than t_a pixels are not ignored. The exact value of t_a can be chosen based on the properties of the document images under consideration (minimum font size, resolution, ...) and the desired geometric accuracy of the evaluation results. For the UW-III document images, we used $t_a = 500$ pixels for zone-level evaluation and $t_a = 100$ pixels for text-line-level evaluation.

If there is more than one significant edge incident to a node in G or in H , the node is considered *oversegmented* or *undersegmented*, respectively. Using these definitions, we can introduce several measures for evaluating a page segmentation algorithm. An illustration of these measures is given in Fig. 3. These measures are defined as follows:

Total correct segmentations (T_c). The total number of one-to-one matches between the ground-truth components and the segmentation components.

Total oversegmentations (T_o). The total number of significant edges that ground-truth components have minus the number of ground-truth components to which at least one significant edge is incident.

Total undersegmentations (T_u). The total number of significant edges that segmentation components have minus the number of segmentation components to which at least one significant edge is incident.

Oversegmented components (C_o). The number of ground-truth components having more than one significant edge.

Undersegmented components (C_u). The number of segmentation components having more than one significant edge.

Missed components (C_m). The number of ground-truth components that did not match any foreground component in the hypothesized segmentation.

False alarms (C_f). Number of components in the hypothesized segmentation that did not match any foreground component in the ground-truth segmentation.

3 ALGORITHMS FOR PAGE SEGMENTATION

We selected six representative algorithms for page segmentation. Furthermore, we have introduced a dummy algorithm to determine a base line of the possible performance. Brief descriptions of each algorithm and its parameters are described in turn in the following.

3.1 Dummy Algorithm

The dummy segmentation algorithm always outputs the whole page as one segment. The purpose of this algorithm is to see how well we can perform without doing anything. Then, the performance of other algorithms can be seen as gains over that achieved by the dummy algorithm. Using the dummy algorithm also highlights limitations of the evaluation scheme, as detailed in Section 4.1.

3.2 X-Y Cut

The x-y cut segmentation algorithm [15], also referred to as the recursive x-y cuts (RXYC) algorithm, is a tree-based top-down algorithm. The root of the tree represents the entire document page. All of the leaf nodes together represent the final segmentation. The RXYC algorithm recursively splits the document into two or more smaller rectangular zones, which represent the nodes of the tree. At each step of the recursion, the horizontal and vertical projection profiles of each node are computed. To compute the valleys in the projection profile histograms, noise removal thresholds t_x^n and t_y^n are used. First, the thresholds t_x^n and t_y^n are scaled linearly based on the current zone's width and height. Then, all bins of the histograms that contain values less than the

scaled thresholds are set to zero. The valleys along the horizontal and vertical directions, v_x and v_y , are then compared to the corresponding predefined thresholds t_x and t_y . If the valley is larger than the threshold, the node is split at the midpoint of the wider of v_x and v_y into two children nodes. The process continues until no leaf node can be split further.

3.3 Smearing

The run-length smearing algorithm (RLSA) [18] works on binary images, where white pixels are represented by 0s and black pixels by 1s. The algorithm transforms a binary sequence x into y according to the following rules:

1. 0s in x are changed to 1s in y if the number of adjacent 0s is less than or equal to a predefined threshold C .
2. 1s in x are unchanged in y .

These steps have the effect of linking together neighboring black areas that are separated by less than C pixels. The RLSA is applied row-wise to the document using a threshold t_{sh} and column-wise using threshold t_{sv} , yielding two distinct bitmaps. These two bitmaps are combined in a logical AND operation. Additional horizontal smearing is done to obtain a smoothed final bitmap using a smaller threshold, t_{sm} . Then, connected component analysis is performed on this bitmap to obtain document zones. The mean horizontal run length R_m of the black pixels in the original image and the mean block height H_m are calculated. Then, a block is classified into a text block if

$$R < f_{tr}R_m \quad \text{and} \quad H < f_{th}H_m, \quad (2)$$

where f_{tr} and f_{th} are two thresholds, R is the horizontal run length of the black pixels in the current block, and H is the block height.

3.4 Whitespace Analysis

The whitespace analysis algorithm described by Baird [19] analyzes the structure of the white background in document images. The first step is to find a set of maximal white rectangles (called *covers*) whose union completely covers the background. Breuel's algorithm for finding the maximal empty whitespace [20] is used in our implementation for this step. These covers are then sorted with respect to the sort key, $K(c)$:

$$K(c) = \sqrt{\text{area}(c) * W(\lceil \log_2(\text{height}(c)/\text{width}(c)) \rceil)}, \quad (3)$$

where c is the cover and $W(\cdot)$ is a dimensionless weighting function. Baird [19] chose a special weighting function using experiments on a particular data set. We used an approximation of the original weighting function as

$$W(x) = \begin{cases} 0.5 & \text{if } x < 3 \\ 1.5 & \text{if } 3 \leq x < 5 \\ 1 & \text{if } x \geq 5. \end{cases} \quad (4)$$

The purpose of the weighting function is to assign higher weight to tall and long rectangles because they are supposed to be meaningful separators of text blocks.

In the second step, the rectangular covers c_i , $i = 1, \dots, m$, where m is the total number of whitespace

covers, are combined one by one to generate a corresponding sequence s_j , $j = 1, \dots, m$ of segmentations. A segmentation is the uncovered area left by the union of the covers combined so far. Before a cover c_i is unified to the segmentation s_j , a trimming rule is applied to avoid early segmentation of narrow blocks. The unification of covers continues until the stopping rule (5) is satisfied:

$$K(s_j) - f_w * j/m \leq t_s, \quad (5)$$

where $K(s_j)$ is the sort key $K(c_j)$ of the last cover unified in making segmentation s_j , f_w is a weighting factor, and t_s is the stopping threshold. At the final segmentation, connected components within the remaining uncovered parts are candidate text regions. Since the uncovered regions thus obtained are not necessarily rectangular in shape, we take bounding boxes of these uncovered regions as representative of the text segments.

3.5 Constrained Textline Detection

The layout analysis approach by Breuel [20] finds text-lines as a three step process:

1. Find empty whitespace rectangles that completely cover the page background. The algorithm for finding maximal empty rectangles is described in [20]. The algorithm returns whitespace rectangles in the order of a decreasing area. The rectangles are allowed a maximum overlap of t_o . Usually, 300 rectangles are sufficient to completely cover the page background.
2. The whitespace rectangles are evaluated as candidates for column separators or gutters based on their aspect ratio, width, and proximity to text-sized connected components.
3. The whitespace rectangles representing the gutters are used as obstacles in a robust least square text-line detection algorithm [32]. Then, the bounding box of all the characters making the text-line is computed.

The method was merely intended by its author as a demonstration of the application of two geometric algorithms and not as a complete layout analysis system; nevertheless, we included it in the comparison because it has already been proven useful in many applications. It is also nearly parameter free and resolution independent.

3.6 Docstrum

The docstrum algorithm proposed by O'Gorman [16] is a bottom-up approach based on the nearest neighborhood clustering of connected components extracted from the document image. After noise removal, the connected components are separated into two groups, one with characters of the dominant font size and another one with characters in titles and section headings, using a character size ratio factor f_d . Then, K nearest neighbors are found for each connected component. A histogram of the distance and angle of each connected component from its K nearest neighbors is computed. The peak of the angle histogram gives the dominant skew in the document image. This skew estimate is used to compute within-line nearest neighbor pairs. Then, text-lines are found by computing the transitive closure on within-line nearest neighbor pairings using a

threshold t_{tc} . Finally, text-lines are merged to form text blocks using a parallel distance threshold t_{pa} and a perpendicular distance threshold t_{pe} .

3.7 Voronoi Diagram-Based Algorithm

The Voronoi diagram-based segmentation algorithm by Kise et al. [17] is also a bottom-up algorithm. In the first step, it extracts sample points from the boundaries of the connected components using a sampling rate r_s . Then, noise removal is done using a maximum noise zone size threshold t_n , in addition to width, height, and aspect ratio thresholds. After that, a Voronoi diagram is generated using sample points obtained from the borders of the connected components. The Voronoi edges that pass through a connected component are deleted to obtain an area Voronoi diagram. Finally, superfluous Voronoi edges are deleted to obtain boundaries of document components. An edge is declared superfluous if it satisfies any of the following criteria:

1. The minimum distance d between its associated connected components is less than the intercharacter gap in body text regions.
2. The minimum distance d between its associated connected components is less than the interline spacing times a margin control factor f_m or the area ratio of the two connected components is above an area ratio threshold t_a .
3. At least one of its terminals is neither shared by another Voronoi edge nor lies on the edge of the document image.

The output of the algorithm consists of arbitrarily shaped regions bounded by Voronoi edges. Since we evaluate all algorithms on document pages with Manhattan layouts, we represent each Voronoi region by its bounding box.

4 EXPERIMENTS AND RESULTS

Based on the performance measures defined in Section 2, we evaluated the performance of six algorithms for page segmentations, namely, x-y cut [15], the smearing algorithm [18], whitespace analysis [19], docstrum [16], the Voronoi diagram-based approach [17], and the constrained text-line finding algorithm [20]. The evaluation of the algorithms was done on the University of Washington III (UW-III) database [21].

The UW-III database consists of 1,600 English document images with Manhattan layouts scanned from different archival journals with manually edited ground truth of entity bounding boxes. These bounding boxes enclose text and nontext zones, text-lines, and words. For each document, a number of page and zone attributes are available as well. The UW-III data set provides a good basis for comparative evaluation of page segmentation algorithms since the majority of documents available today, like books, journals, magazines, letters, etc., have Manhattan layouts. Researchers have used the University of Washington data set for quantitatively evaluating different document analysis tasks like noise removal [33], [34], skew estimation [35], table recognition [24], document zone classification [22], [36], and layout-based document image retrieval [37].

TABLE 1
Parameter Values Used for Each Algorithm
in the Evaluation Given in Table 2

Algorithm	Default values	Optimal values
Dummy	None	
X-Y cut	$t_x = 35, t_y = 54, t_x^n = 78, t_y^n = 32$	
Smearing	$t_{sh} = 300, t_{sv} = 500, t_{sm} = 30, f_{tr} = 3, f_{th} = 3$	
Text-line	$t_o = 0.8$	
Whitespace	$f_w = 42.43, t_s = 34.29$	$f_w = 42.43, t_s = 65$
Docstrum	$K = 5, t_{tc} = 2.578, f_d = 9,$ $t_{pe} = 1.3, t_{pa} = 1.5$	$K = 8, t_{tc} = 2.578, f_d = 9,$ $t_{pe} = 0.6, t_{pa} = 2.345$
Voronoi	$s_r = 6, t_n = 11,$ $f_m = 0.34, t_a = 40$	$s_r = 6, t_n = 11,$ $f_m = 0.083, t_a = 200$

For the dummy, x-y cut, smearing, and text-line finding algorithms, the default and optimized parameters are the same.

Manhattan layouts pose specific constraints on page layout that can be used by page segmentation algorithms to achieve lower error rates. To evaluate the performance of page segmentation algorithms on non-Manhattan layouts, it would be useful to experiment on a data set containing non-Manhattan layouts only. This way, one can find which algorithms are suitable for Manhattan layouts and which algorithms are suitable for non-Manhattan layouts.

We have divided the experiments into two parts:

1. *Benchmarking* of the algorithms based on the text-line-based measure of block segmentation accuracy given by (1).
2. *Performance evaluation* of the algorithms based on the vectorial score defined in Section 2.3.

The first experiment augments the work by Mao and Kanungo [4] and adds three more algorithms to the comparison. We also do a detailed analysis of the errors to show that the limitation of the algorithm as pointed out in Section 2 is reflected in the results. We also show that this gives completely misleading results in certain cases. The second experiment demonstrates the benefits of our vectorial-score-based evaluation method as compared to the single-score-based measure.

4.1 Benchmarking

The benchmarking of the page segmentation algorithms was done on a subset of the UW-III database. We chose the 978 images that correspond to the UW-I data set pages, as was done in [4]. Only the text regions are evaluated and non-text regions are ignored. The data set is divided into 100 training images and 878 test images. The purpose of the training images is to find suitable parameter values for the segmentation algorithms. The experiments are done using both default parameters, as mentioned in the respective papers, and tuned/optimized parameters (Table 1). This allows us to assess how much the performance of each algorithm depends on the choice of good parameters for the task. The parameters for the x-y cut algorithm are highly application dependent, so no default parameters are specified in [15]. The optimized parameter values used for x-y cut, docstrum, and Voronoi-diagram-based algorithms were the same as in [4]. For the smearing, whitespace, and constrained text-line finding algorithms, we experimented with different parameter values and selected those that gave the lowest error rates on the training set.

TABLE 2
The Evaluation Results for Different Page Segmentation Algorithms on 100 Training Images and 878 Test Images

Algorithm	Default parameters			Optimized parameters		
	Train		Test	Train		Test
	Mean	Mean	Stdev	Mean	Mean	Stdev
Dummy	52.2	48.8	39.0	52.2	48.8	39.0
X-Y cut	14.7	17.1	24.4	14.7	17.1	24.4
Smearing	13.4	14.2	23.0	13.4	14.2	23.0
Whitespace	12.7	12.2	20.0	9.1	9.8	18.3
Text-line	8.9	8.5	14.4	8.9	8.5	14.4
Docstrum	8.7	11.2	22.6	4.3	6.0	15.2
Voronoi	6.8	7.5	12.9	4.7	5.5	12.3

The results are reported in terms of percentage of text-lines detection errors (1).

TABLE 3
Text-Line Detection Errors (Percent) for Each of the Algorithms Separated for One, Two, and Three-Column Documents and Separated for Photocopies or Direct Scans

Algorithm	No. of columns			Photocopy	
	1	2	3	No	Yes
Dummy	8.3	75.6	88.5	68.7	46.2
X-Y cut	19.9	15.6	11.7	14.7	17.4
Smearing	23.5	7.9	5.8	6.6	15.1
Whitespace	14.5	6.7	5.6	2.9	10.8
Text-line	13.3	5.3	4.4	3.6	9.2
Docstrum	5.8	6.2	5.2	6.2	5.9
Voronoi	6.9	4.6	3.4	2.8	5.8

We have used the page segmentation evaluation toolkit (PSET) [38] that implements the training and evaluation scheme in [4]. The average text-line detection error rate for each algorithm is given in Table 2. The high standard deviation in the error rate of each algorithm shows that the algorithms work very well on some images, while failing badly on some other images.

Table 3 shows the error rates of the algorithms separated for different document characteristics. First, the documents were separated according to the "maximum columns number" attribute recorded for each page. There are 362, 449, and 67 one, two, and three-column documents in the test set of 878 pages, respectively. We can observe that the smearing, whitespace, and text-line algorithms perform much worse on one-column documents than on average. This behavior can be explained by the stronger effect of the noise blocks occurring in photocopied images for these one-column documents because each line is affected. We further investigated this hypothesis by separating the documents according to their "degradation type" attribute. There are 776 photocopied and 102 directly scanned documents in the test set. The respective results are shown in Table 3. We can observe that the algorithms performing worst on one-column documents in fact also perform worst on the photocopied images due to the noise blocks. Interestingly, the docstrum algorithm especially does not gain accuracy for clean documents, while the Voronoi-based algorithm still performs best. The smearing, whitespace, and text-line algorithms are most affected by the photocopy effects. This suggests that they would perform better for current layout analysis tasks in which most documents are directly scanned.

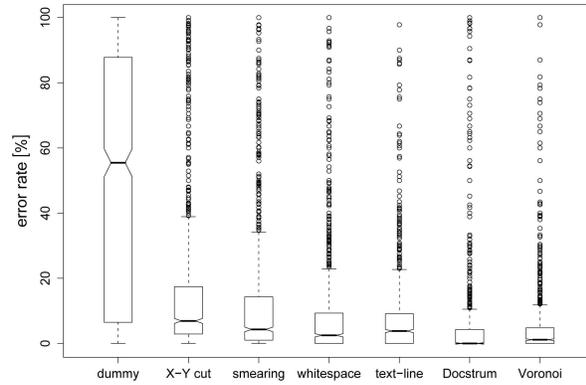


Fig. 4. Box plot for the results obtained with optimized parameters on the test data.

Fig. 4 shows a box plot of the error rates observed for each algorithm. The boxes in the box plot represent the interquartile range, that is, they contain the middle 50 percent of the data. The lower and upper edges represent the first and third quartiles, whereas the middle line represents the median of the data. The notches represent the expected range of the median. The "whiskers" on the two sides show inliers, that is, points within 1.5 times the interquartile range. The outliers are represented by small circles outside the whiskers. We can observe the following details: A ranking of the algorithms based on their median error would deviate from the ranking based on the average error. Remarkably, the docstrum algorithm does not make any errors for more than 50 percent of the pages in the test set. This performance is not achieved by any other algorithm. This might be a property that would be preferable in certain applications, while, for other applications, the average error rate may be more important.

To study the similarities in the behavior of different algorithms, we plot the correlation of the errors made by each algorithm in Fig. 5. Each dot in the correlation plot represents one document image. The horizontal and vertical axis represent the error made by the corresponding algorithms. It can be seen from the correlation plot that the docstrum and Voronoi algorithms show strong correlation because they both are bottom-up approaches. Also, the x-y cut and the dummy algorithm are highly correlated. This is due to the fact that the x-y cut algorithm fails on documents with a large amount of noise and reports the whole page as one segment, which is the same output as generated by the dummy algorithm. When this happens for a single column document, the error rate computed by (1) is zero. However, in the case of single-column documents with a large amount of noise, it is not possible to segment them into text-lines merely by horizontal projection. Hence, the error rates reported in these cases give misleading results. An example of such a document from the test set is shown in Fig. 6. Since there are only a few images in the test set that fall into this category, the experimental results are still valid. An interesting observation that can be made from the correlation plot is that, for each algorithm, there are some documents on which it performs better than all of the other algorithms. This indicates that combining the output of more than one algorithm might yield better results.

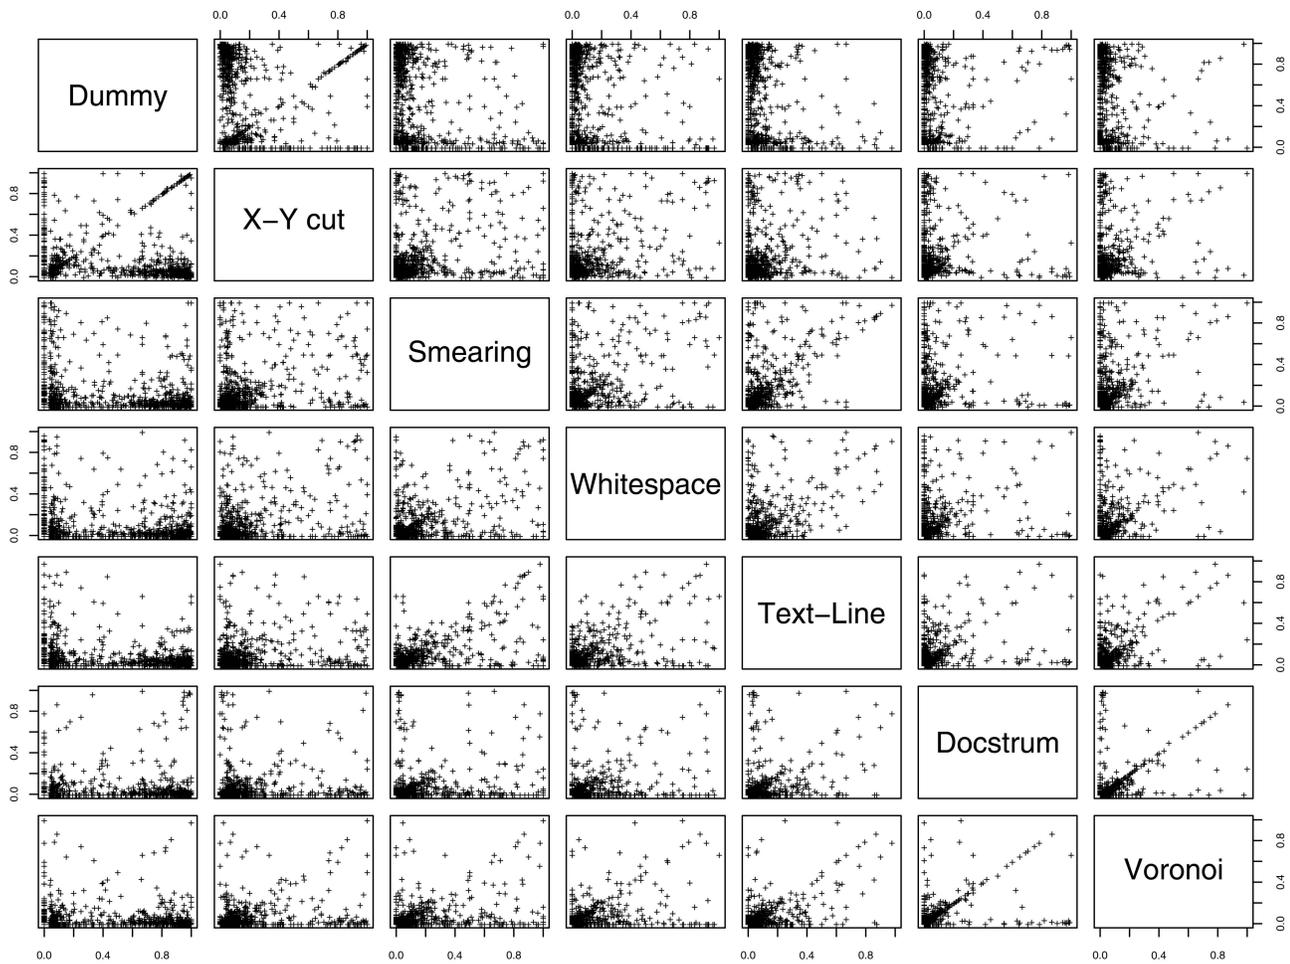


Fig. 5. Correlation plot of the errors made by each algorithm. Each dot in the correlation plot represents one document image. The horizontal and vertical axes represent the error made by the corresponding algorithms.

4.2 Performance Evaluation

The performance of the six-page segmentation algorithms was evaluated on the complete UW-III data set based on the measures defined in Section 2.3. These measures evaluate different aspects of a page segmentation algorithm for a given parameter setting. The goal of these performance measures is not to optimize the parameters of an algorithm on this basis because the importance of different measures is entirely application-dependent. If an OCR system expects single text-line images as input, undersegmentation (for example, putting two consecutive lines together) poses a much more serious problem than oversegmentation (like segmenting a text-line into words). If the OCR system accepts both text-lines and text blocks as input, the only major problem is undersegmentation (for example, merging two text columns). In any case, the ground truth should also fulfill the demands of the target application. For instance, for single-line OCR, text-line-level ground truth should be used, although, for block-level OCR, either text-column or text-zone level ground truth should be used. Since the parameters of the page segmentation algorithms given in Table 1 were optimized with respect to block-level OCR application, these parameters can be used in these evaluations as well. The parameters for the x-y cut, whitespace analysis, docstrum, and Voronoi diagram-based algorithms were tuned to segment text zones. Hence, they were evaluated on zone-level ground truth with the results given

in Table 4. The smearing and the constrained text-line finding algorithms locate text-lines in the given image. Therefore, they are evaluated on text-line-level ground truth with the results given in Table 5.

A problem with the text-zone-level ground truth, in the UW-III data set, is that a single paragraph is considered one text zone. Hence, two consecutive paragraphs on the same page make two different zones. In many documents, the segmentation of text columns into paragraphs is indicated by indentation rather than spacing. Determining paragraphs from indentations is usually a separate processing step. Therefore, an evaluation based on paragraph-level ground truth may not correctly reflect the performance of a page segmentation algorithm by giving more undersegmentation errors than the algorithm actually made.

We modified the ground truth for UW-III to get text zones instead of paragraphs. For this purpose, we first specified a partial order of the text paragraphs based on their spatial relationships and then used a topological sorting algorithm to find the reading order, as in [23]. Then, the bounding boxes of two consecutive paragraphs in the reading order were merged if their start and end positions along the horizontal direction are within five pixels of each other. These modified text zones were used to evaluate the page segmentation algorithms with the results, as shown in Table 6.

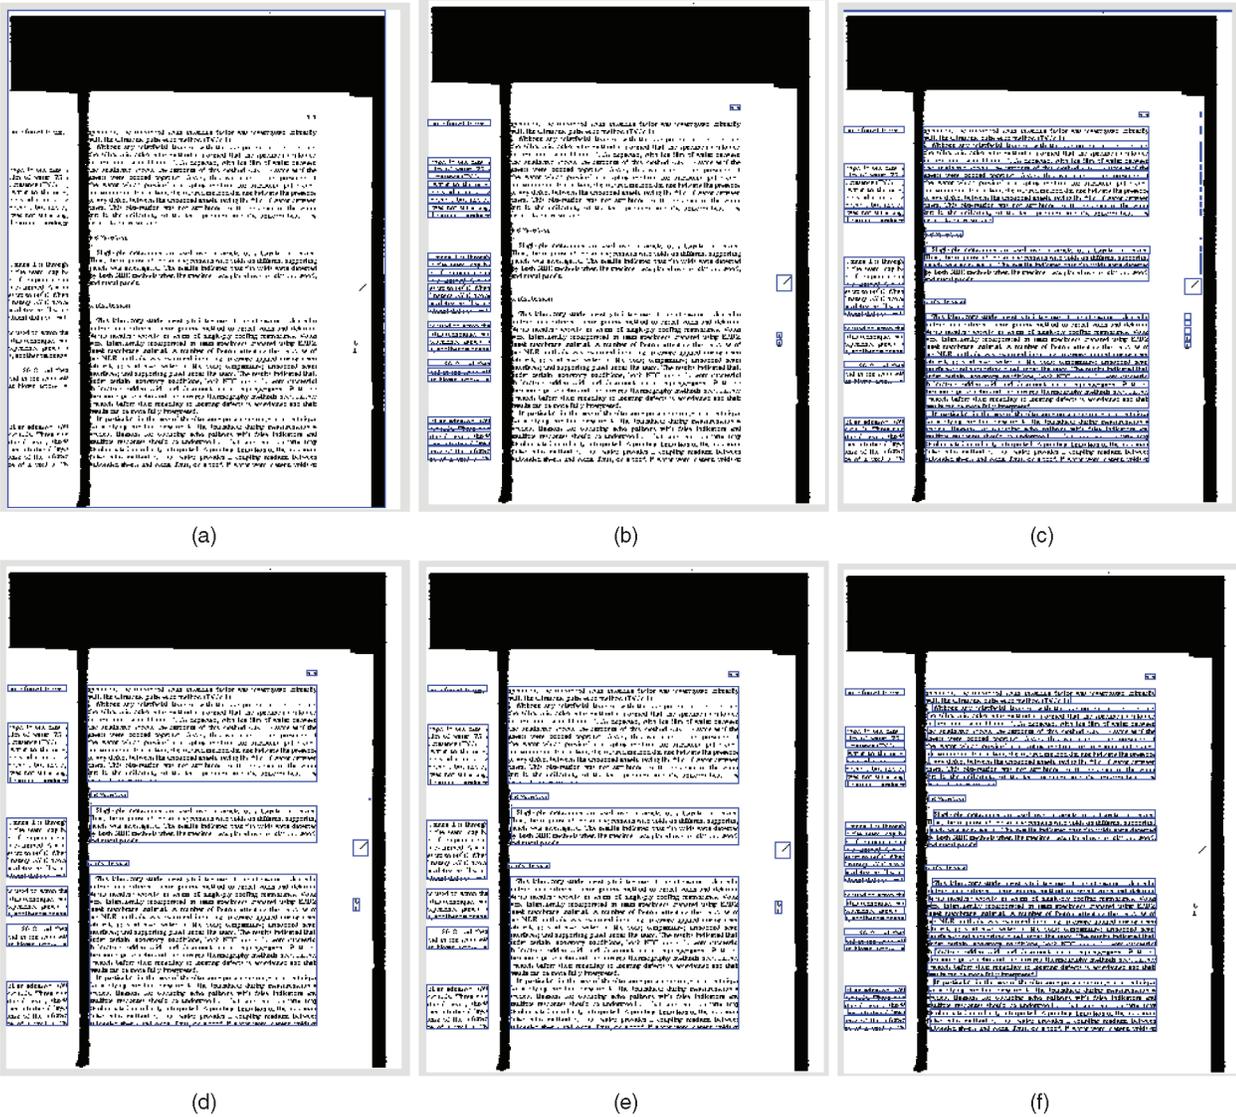


Fig. 6. Segmentation results from applying each algorithm to an image (D047) in the test set. The error rates calculated according to [4] (1) show that the x-y cut algorithm performs the best in this case, which is clearly misleading. (a) X-Y cut ($\rho = 0.000$). (b) Smearing ($\rho = 0.976$). (c) Whitespace ($\rho = 0.463$). (d) Docstrum ($\rho = 0.561$). (e) Voronoi ($\rho = 0.561$). (f) Text-line ($\rho = 0.756$).

TABLE 4
Different Types of Errors Made by Each Algorithm on the Original Zone-Level Ground Truth

Algorithm	Segmented zones	T_c	T_o	T_u	C_o	C_u	C_m	C_f
Dummy	6.60	0.00	0.00	93.34	0.00	6.60	0.00	0.00
X-Y cut	61.74	19.66	28.29	57.23	11.94	17.70	1.73	24.90
Whitespace	84.27	35.22	28.78	43.29	11.98	18.27	0.47	31.10
Docstrum	155.88	44.13	81.23	30.38	13.26	13.94	1.34	67.89
Voronoi	165.22	38.34	92.32	34.59	14.57	15.25	1.72	42.21

Each text paragraph is considered a separate text zone. All entries are normalized by the total number of zones, 24,247, and are expressed as percentage. The column labels are total correct segmentations (T_c), total oversegmentations (T_o), total undersegmentations (T_u), oversegmented components (C_o), undersegmented components (C_u), missed components (C_m), and false alarms (C_f).

TABLE 5
Different Types of Errors Made by Each Algorithm on Text-Line-Level Ground Truth

Algorithm	Segmented lines	T_c	T_o	T_u	C_o	C_u	C_m	C_f
Textline	100.13	97.17	3.77	1.64	2.82	1.23	0.26	36.05
Smearing	98.55	92.82	3.30	1.25	1.64	0.86	3.92	38.24

For a key to column labels, please refer to Table 4. All entries are normalized by the total number of text-lines, 105,443, and are expressed as percentage.

TABLE 6
Different Types of Errors Made by Each Algorithm on Modified Zone-Level Ground Truth

Algorithm	Segmented zones	T_c	T_o	T_u	C_o	C_u	C_m	C_f
Dummy	8.05	0.00	0.00	91.88	0.00	8.04	0.00	0.00
X-Y cut	74.55	22.07	37.18	51.83	16.29	19.32	1.95	31.09
Whitespace	101.61	37.41	40.18	37.15	16.49	18.45	0.52	39.05
Docstrum	188.68	45.73	105.99	23.54	21.63	13.43	1.23	84.16
Voronoi	199.89	40.68	118.18	27.43	22.02	14.96	1.66	53.03

For a key to column labels, please refer to Table 4.



Fig. 7. Segmentation results from applying each algorithm to an image (A005) in the test set. The page contains a title in large font and a big noise strip along the right border. (a) The x-y cut algorithm fails in the presence of noise and tends to take the whole page as one segment. (b) The smearing algorithm also classifies the detected regions as text/nontext and thus misses the lines joined by the noise bar. The (c) whitespace, (d) docstrum, and (e) Voronoi algorithms split the title lines due to the large font size and big interword spacing. (f) The text-line finding algorithm. Due to the noise bar, several characters on the right side of each line in the second column were merged with the noise bar and the text-line finding algorithm did not include these characters.

The results of applying each algorithm to an example image are shown in Fig. 7. Based on the results in Tables 5 and 6, we can make the following observations about each algorithm:

- The dummy algorithm has no correct segmentations, and all of the components are undersegmented.
- The x-y cut algorithm fails in the presence of noise and tends to take the whole page as one segment. This results in many undersegmentation errors.
- The whitespace algorithm is sensitive to the stopping rule. Early stopping results in a higher number of undersegmentation errors, while late stopping results in more oversegmentation errors. The whitespace

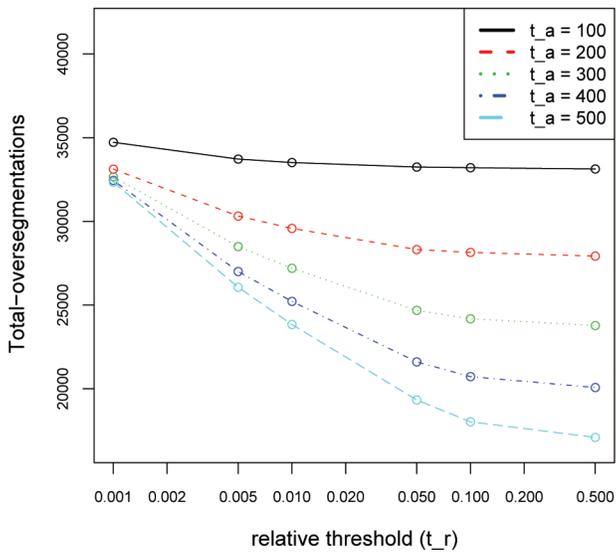


Fig. 8. A plot of the values of total oversegmentations made by the Voronoi algorithm as the values of thresholds t_r and t_a defining significant edges are changed.

algorithm also made few missed errors because all connected components with width larger than half the page width or height greater than half the page height were removed prior to the computation of white-spaces. Hence, separator lines in the header or footer, which are considered as zones in UW-III ground truth, were missed by the algorithm.

- In the Voronoi and docstrum algorithms, the inter-character and interline spacings are estimated from the document image. Hence, spacing variations due to different font sizes and styles within one page result in oversegmentation errors in both algorithms. For instance, in many cases, they fail to estimate the interline distance correctly and, hence, split the zones into individual text-lines, resulting in a large number of oversegmentation errors. The number of segmented zones for these two algorithms is much higher than the number of zones in the ground truth. In some cases, text-lines in page titles are incorrectly segmented (see Fig. 7) due to a large variation in font size.
- The smearing algorithm classifies text-lines merged with noise blocks as nontext, resulting in a large number of missed errors.
- The major part of the errors made by the constrained text-line finding algorithm are missed errors. Single digit page numbers are missed by the text-line finding algorithm because it requires at least two connected components to form a line. In some cases, the characters from two consecutive lines are merged. Hence, the bounding box of the lower text-line spans across both text-lines, resulting in both oversegmentation and undersegmentation errors.

The choice of the values of thresholds t_r and t_a defining significant edges is application-dependent. In the case of OCR, it might be important to keep the thresholds low so that even a missed dot is reported as an error. However, other applications, like layout-based document image retrieval, have less strict demands on the geometric accuracy

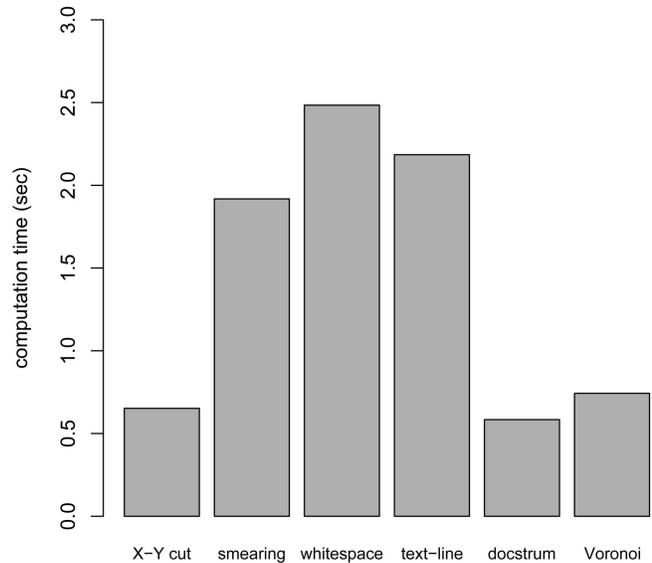


Fig. 9. Average runtime for each algorithm on the UW-III data set. The experiment was run on an AMD Opteron 2.4 GHz machine running Linux.

of page segmentation. To evaluate the sensitivity of the performance measures with respect to the thresholds t_r and t_a , we have conducted an experiment. We have chosen the Voronoi algorithm as a sample page segmentation algorithm and have observed the changes in the number of reported total oversegmentation errors as the values of the thresholds t_r and t_a are varied over a broad range. The algorithm was run over the complete UW-III data set. Then, the output was compared to the zone-level ground truth using different combinations of t_r and t_a . The resulting plot is shown in Fig. 8. From the plot, it can be noticed that setting either t_r or t_a to a very low value makes the performance measure independent of the other threshold. As expected the number of detected total oversegmentations decreases when the values of both thresholds are increased simultaneously. For OCR applications, just setting t_a to a very small value, for instance, equal to the size of a dot, and ignoring t_r altogether might be a good choice. In the case of layout-based retrieval, we have to consider both thresholds because the size of small zones like page numbers might be smaller than a moderately chosen value of t_a . In such a case, t_r helps by keeping the threshold low for small zones.

The average runtime of the evaluated page segmentation algorithms is shown in Fig. 9. The timing of the algorithms cannot be directly compared because of the differences in their input and output. The whitespace, docstrum, Voronoi, and x-y cut algorithms give text blocks that still have to be separated into text-lines, whereas the constrained text-line finding algorithm directly gives the text-lines as output. Second, the smearing algorithm also includes a block-classification step, which is missing in other algorithms. Furthermore, the docstrum, whitespace, and constrained text-line finding algorithms depend on the computation of connected components in the image, which were calculated offline and stored in the database. In general, the x-y cut, docstrum, and Voronoi algorithms took less than half the time as compared to the smearing, whitespace analysis, and constrained text-line finding algorithms.

4.3 Recommendations

Based on the experimental results and observations, the following recommendations can be made about the choice of page segmentation algorithm for different applications and document types.

- For clean documents with little or no skew, the x-y cut algorithm might be a good choice as it is fast and easy to implement.
- For a homogeneous collection of documents (same resolution, similar layouts, similar font sizes, and styles) with a variable amount of noise, the docstrum and Voronoi algorithms can be used. However, the parameters of these algorithms should be tuned to segment the given document collection to obtain good results.
- For documents containing many font sizes and styles, the constrained text-line finding algorithm works best because it is based on geometric models that are invariant to font size, font style, and scan resolution.
- For a diverse document collection with documents having different font sizes and layouts or documents scanned at different resolutions, the constrained text-line finding algorithm is a good choice because it is nearly parameter-free.
- For non-Manhattan layouts, or layouts having text in different orientations, the Voronoi algorithm can be a good choice.

One problem with the evaluated page segmentation algorithms is that they give a single segmentation of a page without any confidence value. Therefore, if the output has to be verified manually, one has to look at the segmentation done for each page individually. This can be very cumbersome, even prohibitive, for large-scale applications like Google book search [39]. One solution to this problem is to do page segmentation in a probabilistic framework, allowing the operator to look at only those pages for which the confidence value returned by the algorithm is low. Hence, in our opinion, an important direction of future research in page segmentation will be to develop probabilistic algorithms that can handle real-world documents.

5 CONCLUSION

We presented an approach for evaluating page-segmentation algorithms using color-based representation. The color-based representation of segmentation is independent of zone shape and it can be saved and exchanged using any lossless color image format. Instead of using a single score for the performance of each algorithm, different aspects of the algorithms are evaluated separately. Depending on the target application, different error measures may be weighted according to their significance in that application. Using these performance measures, we have analyzed the strengths and weaknesses of six popular algorithms for page segmentation.

Our experiments showed that the x-y cut and the smearing algorithms fail to segment a page in the presence of noise. The whitespace analysis algorithm is sensitive to the stopping rule and results in either oversegmentations or

undersegmentations. The docstrum and the Voronoi algorithms tend to oversegment title and section headings if the font size is much different from body text in that page. The constrained text-line finding algorithm misses single-digit page numbers as it requires at least two components to make a line.

Based on our experiments, we can conclude that, for a homogeneous document collection with a large proportion of documents with Manhattan layouts, the docstrum and Voronoi algorithms are the best choice. In the case of a heterogeneous document collection with different font sizes, styles, and scan resolutions, the constrained text-line finding algorithm appears to be the best choice.

ACKNOWLEDGMENTS

This work was partially funded by the BMBF (German Federal Ministry of Education and Research), project IPeT (01 IW D03).

REFERENCES

- [1] R. Cattoni, T. Coianiz, S. Messelodi, and C.M. Modena, "Geometric Layout Analysis Techniques for Document Image Understanding: A Review," IRST Technical Report 9703-09, 1998.
- [2] G. Nagy, "Twenty Years of Document Image Analysis in PAMI," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 1, pp. 38-62, Jan. 2000.
- [3] S. Mao, A. Rosenfeld, and T. Kanungo, "Document Structure Analysis Algorithms: A Literature Survey," *Proc. SPIE Electronic Imaging*, vol. 5010, pp. 197-207, Jan. 2003.
- [4] S. Mao and T. Kanungo, "Empirical Performance Evaluation Methodology and Its Application to Page Segmentation Algorithms," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 3, pp. 242-256, Mar. 2001.
- [5] F. Lotti, P. Heroux, S. Adam, G. Sanchez, E. Valveny, P. Dosch, and J. Lladós, "Performance Analysis and Evaluation Working Group Report," *Document Analysis Systems*, <http://www.dsi.unifi.it/DAS04/DASPerfEv.pdf>, Sept. 2004.
- [6] A. Antonacopoulos, D. Karatzas, and D. Bridson, "Ground Truth for Layout Analysis Performance Evaluation," *Document Analysis Systems*, pp. 302-311, Feb. 2006.
- [7] A. Antonacopoulos, B. Gatos, and D. Karatzas, "ICDAR 2003 Page Segmentation Competition," *Proc. Seventh Int'l Conf. Document Analysis and Recognition*, pp. 688-692, 2003.
- [8] A. Antonacopoulos, B. Gatos, and D. Bridson, "ICDAR 2005 Page Segmentation Competition," *Proc. Eighth Int'l Conf. Document Analysis and Recognition*, pp. 75-80, Aug. 2005.
- [9] J. Kanai, T.A. Nartker, S.V. Rice, and G. Nagy, "Performance Metrics for Document Understanding Systems," *Proc. Second Int'l Conf. Document Analysis and Recognition*, pp. 424-427, Oct. 1993.
- [10] B.A. Yanikoglu and L. Vincent, "Ground-Truthing and Benchmarking Document Page Segmentation," *Proc. Third Int'l Conf. Document Analysis and Recognition*, pp. 601-604, Aug. 1995.
- [11] J. Liang, I.T. Phillips, and R.M. Haralick, "Performance Evaluation of Document Structure Extraction Algorithms," *Computer Vision and Image Understanding*, vol. 84, pp. 144-159, 2001.
- [12] A.K. Das, S.K. Saha, and B. Chanda, "An Empirical Measure of the Performance of a Document Image Segmentation Algorithm," *Int'l J. Document Analysis and Recognition*, vol. 4, no. 3, pp. 183-190, 2002.
- [13] A. Hoover, G. Jean-Baptiste, X. Jiang, P.J. Flynn, H. Bunke, D.B. Goldgof, K. Bowyer, D.W. Eggert, A. Fitzgibbon, and R.B. Fisher, "An Experimental Comparison of Range Image Segmentation Algorithms," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 7, pp. 673-689, July 1996.
- [14] X. Jiang, C. Marti, C. Irniger, and H. Bunke, "Distance Measures for Image Segmentation Evaluation," *EURASIP J. Applied Signal Processing*, vol. 2006, Article ID 35 909, 2006.
- [15] G. Nagy, S. Seth, and M. Viswanathan, "A Prototype Document Image Analysis System for Technical Journals," *Computer*, vol. 25, no. 7, pp. 10-22, July 1992.

- [16] L. O'Gorman, "The Document Spectrum for Page Layout Analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 11, pp. 1162-1173, Nov. 1993.
- [17] K. Kise, A. Sato, and M. Iwata, "Segmentation of Page Images Using the Area Voronoi Diagram," *Computer Vision and Image Understanding*, vol. 70, no. 3, pp. 370-382, June 1998.
- [18] K.Y. Wong, R.G. Casey, and F.M. Wahl, "Document Analysis System," *IBM J. Research and Development*, vol. 26, no. 6, pp. 647-656, 1982.
- [19] H.S. Baird, "Background Structure in Document Images," *Document Image Analysis*, H. Bunke, P. Wang, and H.S. Baird, eds., pp. 17-34, World Scientific, 1994.
- [20] T.M. Breuel, "Two Geometric Algorithms for Layout Analysis," *Document Analysis Systems*, pp. 188-199, Aug. 2002.
- [21] I. Guyon, R.M. Haralick, J.J. Hull, and I.T. Phillips, "Data Sets for OCR and Document Image Understanding Research," *Handbook of Character Recognition and Document Image Analysis*, H. Bunke and P. Wang, eds., pp. 779-799, World Scientific, 1997.
- [22] Y. Wang, R. Haralick, and I. Phillips, "Document Zone Content Classification and Its Performance Evaluation," *Pattern Recognition*, vol. 39, no. 1, pp. 57-73, Jan. 2006.
- [23] T.M. Breuel, "High Performance Document Layout Analysis," *Proc. Symp. Document Image Understanding Technology*, Apr. 2003.
- [24] S. Mandal, S. Chowdhury, A. Das, and B. Chanda, "A Simple and Effective Table Detection System from Document Images," *Int'l J. Document Analysis and Recognition*, vol. 8, nos. 2-3, pp. 172-182, June 2006.
- [25] C. Shin and D. Doermann, "Classification of Document Page Images," *Proc. Symp. Document Image Understanding Technology*, pp. 166-175, Apr. 1999.
- [26] F. Shafait, D. Keysers, and T.M. Breuel, "Performance Comparison of Six Algorithms for Page Segmentation," *Proc. Seventh IAPR Workshop Document Analysis Systems*, pp. 368-379, Feb. 2006.
- [27] F. Shafait, D. Keysers, and T.M. Breuel, "Pixel-Accurate Representation and Evaluation of Page Segmentation in Document Images," *Proc. 18th Int'l Conf. Pattern Recognition*, pp. 872-875, Aug. 2006.
- [28] D. Dori, D. Doermann, C. Shin, R. Haralick, I. Phillips, M. Buchman, and D. Ross, "The Representation of Document Structure: A Generic Object-Process Analysis," *Handbook of Character Recognition and Document Image Analysis*, H. Bunke and P. Wang, eds., pp. 421-456, World Scientific, 1997.
- [29] G. Ford and D. Thoma, "Ground Truth Data for Document Image Analysis," *Proc. Symp. Document Image Understanding and Technology*, pp. 199-205, Apr. 2003.
- [30] F. Shafait and T.M. Breuel, "Document Image Dewarping Contest," *Proc. Second Int'l Workshop Camera-Based Document Analysis and Recognition*, pp. 181-188, Sept. 2007.
- [31] T.M. Breuel, "Representations and Metrics for Off-Line Handwriting Segmentation," *Proc. Eighth Int'l Workshop Frontiers in Handwriting Recognition*, pp. 428-433, Aug. 2002.
- [32] T.M. Breuel, "Robust Least Square Baseline Finding Using a Branch and Bound Algorithm," *Proc. Document Recognition and Retrieval VIII*, 2002.
- [33] L. Cinque, S. Levaldi, L. Lombardi, and S. Tanimoto, "Segmentation of Page Images Having Artifacts of Photocopying and Scanning," *Pattern Recognition*, vol. 35, pp. 1167-1177, 2002.
- [34] F. Shafait, J. van Beusekom, D. Keysers, and T.M. Breuel, "Page Frame Detection for Marginal Noise Removal from Scanned Documents," *Proc. 15th Scandinavian Conf. Image Analysis*, pp. 651-660, June 2007.
- [35] O. Okun, M. Pietikainen, and J. Sauvola, "Robust Skew Estimation on Low-Resolution Document Images," *Proc. Fifth Int'l Conf. Document Analysis and Recognition*, pp. 621-624, Sept. 1999.
- [36] D. Keysers, F. Shafait, and T.M. Breuel, "Document Image Zone Classification—A Simple High-Performance Approach," *Proc. Second Int'l Conf. Computer Vision Theory and Applications*, pp. 44-51, Mar. 2007.
- [37] S. Marinai, E. Marino, and G. Soda, "Layout Based Document Image Retrieval by Means of XY Tree Reduction," *Proc. Eighth Int'l Conf. Document Analysis and Recognition*, pp. 432-436, Aug. 2005.
- [38] S. Mao and T. Kanungo, "Software Architecture of PSET: A Page Segmentation Evaluation Toolkit," *Int'l J. Document Analysis and Recognition*, vol. 4, no. 3, pp. 205-217, 2002.
- [39] L. Vincent, "Google Book Search: Document Understanding on a Massive Scale," *Proc. Ninth Int'l Conf. Document Analysis and Recognition*, pp. 819-823, Sept. 2007.



Faisal Shafait received the bachelor's degree in electrical engineering from the University of Engineering and Technology, Taxila, Pakistan, and the master's degree in information and communication systems from the Hamburg University of Technology, Germany. He is currently working toward the PhD degree in the Computer Science Department at the Technical University of Kaiserslautern. He is also a researcher in the Image Understanding and



Pattern Recognition Research Group, German Research Center for Artificial Intelligence, Kaiserslautern, Germany. His research interests include image processing and pattern recognition, focusing on document image analysis.

Daniel Keysers received the Dipl. degree (with honors) and the PhD degree (summa cum laude) in computer science from RWTH Aachen University, Germany, in 2000 and 2006, respectively. During his PhD studies, he was with the Department of Computer Science at RWTH and headed the Image Processing and Understanding Group, Human Language Technology and Pattern Recognition Chair. He visited the Instituto Tecnológico de Informática, Universidad



Politécnica de Valencia, Spain, in 2002 and Microsoft Live Labs in 2006. From 2005 to 2007, he was a senior researcher at the German Research Center for Artificial Intelligence (DFKI), Image Understanding and Pattern Recognition Group (IUPR), Kaiserslautern, Germany. He is currently working at Google Switzerland. His research interests include pattern recognition and statistical modeling, especially for computer vision, image object recognition, image retrieval, and document processing.

Thomas M. Breuel received degrees from the Massachusetts Institute of Technology and Harvard University. Previously, he was a researcher at Xerox PARC, IBM Almaden Research Center, and IDIAP, Switzerland, as well as a consultant to the US Bureau of the Census. He is a professor of computer science in the Computer Science Department at the Technical University of Kaiserslautern, Germany, head of the Image Understanding and Pattern Recognition (IUPR) Research Group at DFKI, and a consultant in Palo Alto, California. His research group works on image understanding, document imaging, computer vision, and pattern recognition.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.