

Document Image Applications

Dan S. Bloomberg and Luc Vincent

Google

Draft for chapter in Livre Hermes Morphologie Mathématique: July 2007

1 Introduction

The analysis of document images is a difficult and ill-defined task. Unlike the graphics operation of rendering a document into a pixmap, using a structured page-level description such as pdf, the analysis problem starts with the pixmap and attempts to generate a structured description. This description is hierarchical, and typically consists of two interleaved trees, one giving the physical *layout* of the elements and the other affixing *semantic* tags. Tag assignment is ambiguous unless the rules determining structure and rendering are tightly constrained and known in advance.

Although the graphical rendering process invariably loses structural information, much useful information can be extracted from the pixmaps. Some of that information, such as skew, warp and text orientation detection, is related to the digitization process and is useful for improving the rendering on a screen or paper. The layout hierarchy can be used to reflow the text for small displays or magnified printing. Other information is useful for organizing the information in an index, or for compressing the image data. This chapter is concerned with robust and efficient methods for extracting such useful data.

1.1 Viewpoint

We start with an empirical observation: *a very large set of document image analysis (DIA) problems can be accurately and efficiently addressed with image morphology and related image processing methods.* Other tools and representations, such as those used in computational geometry, can be quite useful, but they are not required for the vast majority of applications.

We take the view that image analysis is a nonlinear decision process, and that image processing operations are ubiquitously useful. Consequently, we attempt to make decisions based on nonlinear image operations. Many benefits accrue from using the image as the fundamental representation: (1) analysis is very fast; (2) analysis retains the image geometry, so that processing errors are obvious, the accuracy of results is visually evident, and the operations are easily improved; (3) alignment between different renderings and resolutions is maintained; (4) pixel labelling is made in parallel by neighbors; (5) sequential (e.g., filling) operations are used where pixels can have arbitrarily long-range effects, in analogy to IIR filters; (6) pixel groupings are easily determined; (7) segmentation output is naturally represented using masks; (8) implementation is simplified because a relatively small number of operations must be implemented efficiently, and operations on alternate representations are avoided; (9) applications can use both shape

Examples	Constraint	Approach
letterforms	high	Bayesian MAP
page layout	moderate	morphology with params
natural scenes	low	ad hoc

Table 1: Effect of constraints on the approach to image analysis

and texture, at multiple resolutions, to label pixels; and (10) the statistical properties of pixels and sets of pixels can be used to make robust estimation.

With this approach, some operations become trivial. For example, to extract the words from a scanned 1 *bpp* (bit/pixel) image of a music score, a large horizontal morphological erosion generates "seeds" in the staff lines. Then a binary reconstruction (seed fill), using the original image as a mask, recovers the lines and everything touching them. Lyrics and other musical notations are then extracted by XORing with the original.

Table 1 depicts *document image analysis* (DIA) as occupying a high to intermediate position terms of constraints, which depend on the accuracy of the statistical models representing the collection of images. Bayesian statistical models are the most constrained. Analysis is performed by generation from the models, using *maximum a posteriori* (MAP) inference. These techniques have been used for OCR [10] and for locating textlines [9], and can be implemented efficiently using heuristics despite the fact that they require matching all templates at all possible locations [12].

Many DIA problems are not framed in a strict Bayesian format. Although the models are not well-specified, there exist regularities that allow identification of layout parameters (such as average spacing between words and text lines) and, eventually, the layout hierarchy itself. This involves use of both shape and texture, for which morphological operations are ideally suited. At the other extreme, arbitrary natural scenes have very few constraints and continue to defy general attempts at analysis.

In addition to using nonlinear imaging operations to make decisions, it is important to perform operations at scales appropriate for the components under investigation, and 1 *bpp* images typically suffice. Shape and texture play major roles. As the resolution is decreased, component shape is transformed into the texture of page elements farther up the hierarchy. For example, suppose the objective is to identify lines of text. At high resolution, text is composed of letterforms, and one can locate each connected component on the page and then attempt to find textlines by merging their bounding boxes. This is fragile because textlines can be connected by *foreground* (fg) noise. A much simpler method is to reduce the resolution in such a way that the textlines become, in effect, solid horizontal lines. These can be distinguished from halftone images by the fact that the lines of text have white space between them, so that a vertical opening of sufficient size will remove them. We explore page segmentation in more detail later.

1.2 Tools

Here we describe some of the most useful image processing tools for DIA.

The most important low-level operations for DIA fall into five classes:

- *Morphological*. Operations on binary images are by far the most common.

- *Rasterop*. Ubiquitous bit-level operations, these are used for implementing binary morphology, binary logic (e.g., painting and masking) over arbitrary rectangles.
- *Rank reduction*. Nonlinear operations where the subsampled dest pixels are determined using a rank threshold on a tile of pixels from the src, both for binary and grayscale images.
- *Binary reconstruction*. Operations that fill into a mask image from a seed image. These are crucial for accurate segmentation.
- *Connected components*. This differs from the first four operations in that it reads and writes single pixels rather than full words, and can generate non-image data, such as bounding boxes.

These operations can all be implemented efficiently. The first three are parallel: each dest pixel depends only on src pixels. The last two should be done sequentially: the order of operations matters because each dest pixel can depend on previously computed dest pixels. Sequential operations allow a src pixel to affect a dest pixel an arbitrary distance away, whereas parallel operations have a limited extension of influence.

There are two efficient methods for implementing binary morphology: (1) full image rasterop, where each hit or miss in the *structuring element* (Sel) requires a rasterop of src with dest, and (2) destination word accumulation (DWA), where each dest word is computed sequentially using the entire Sel. The DWA method is typically about four times faster than using full image rasterops, due to loop unrolling, fewer cache misses, and fewer writes to memory.

In practice, by far the most common morphological operations use brick Sels, which are separable. For every two-fold reduction in resolution, morphological operations increase in speed by about a factor of 8, where 4x of the speedup is due to the size of the image and the rest comes from using smaller Sels. This forms the basis of the rule that *operations should take place at the lowest resolution that gives the desired accuracy*.

When operations at high resolution are required, the Sels can be very large, and it is important to decompose them into a multiresolution sequence of combs. Most of the gain occurs in going from one to two levels. For example, without decomposition, a horizontal brick Sel of size 64×1 takes 64 rasterops to implement. With two-level decomposition, only $2 * \sqrt{64} = 16$ rasterops are required. This compares well with maximal decomposition of $2 * \log 64 = 12$, so that as a practical matter, two levels of decomposition are sufficient.

The $2x$ binary rank reduction operation can be implemented very efficiently for all four rank levels [1]. For the special case where the decision is based on whether or not at least one pixel in the tile is fg (rank level 1, the max), rank reduction is equivalent to a $2x2$ dilation followed by $2x$ subsampling. Likewise, for the case where all pixels are required to be fg (rank level 4, the min), rank reduction is equivalent to $2x2$ erosion plus subsampling. Similar rank order operations over NxN tiles on grayscale images can be used to avoid the cost of grayscale dilation or erosion at higher resolution followed by subsampling. These are simply computed by finding the max or min, respectively, over each tile, and saving them in a NxN reduced result.

Programs that generate the output shown in the following applications are indicated in the captions. Source code for many of the algorithms described here, including all the examples, can be obtained at www.leptonica.org/source/livre_dia.tar.gz.

2 Applications

We have space to demonstrate a small number of the document image applications that benefit by using a morphological approach.

2.1 Page segmentation

Segmentation is the fundamental operation in DIA. There are many variations and approaches, depending on the goals of the analysis. The goals can be partially specified by the pixel accuracy desired and the cost of various errors. Examples of such goals are:

- Is there an image (or textblock) on the page?
- If there are images (or textblocks), where are they?
- Are there other graphics elements on the page?
- Locate the hierarchical (tree) structure of the text: blocks, paragraphs, sentences, words, characters.
- Assign logical labels to page elements

For a real application, the situation is more nuanced. For example, if the primary goal is good visual appearance, and the non-image part is quantized into a small number of levels, the cost of identifying image pixels as non-image can be much higher than making the opposite mistake. By contrast, if the goal were to identify all the text as a preprocessing step for OCR, it is much worse to lose text regions than to label some image pixels as non-image.

It is often useful to express the page elements as a series of binary masks. Each pixel in a binary mask represents a yes/no decision about whether that pixel has a particular label. Pixels can be represented as fg in multiple masks, such as a pixel that is labeled as fg in both a textline mask and a textblock mask. For example, a halftone mask, with fg pixels over pixels in halftone regions, can be used to remove those image pixels before doing text analysis, or to direct an operation to render the image and non-image pixels differently. The latter is often desirable because text is best rendered with high contrast, whereas images are usually rendered with dithering on printers or with many levels on displays to avoid posterization. In the following, we show how to start with an image and progressively filter different regions, using the implicit shape and texture properties.

Let us first show the use of rank reduction to answer the question: Is there an image on the page? Figure 1 shows the sequence of images. Although a sequence of reductions is taking place, the results are all displayed at the same resolution. Starting with a 300 ppi image containing 8×10^6 pixels (a), do a cascade of four 2x rank reductions. Parts (b) and (c) show the results at 4x and 16x reduction, using levels 1 and 4 followed by 4 and 3, resp. A final 5×5 erosion yields the result (d), and a test for fg pixels gives the answer. This is a computationally inexpensive procedure, taking only 1 msec on a standard 3 GHz processor! This result can be used as a seed in a binary reconstruction to generate the halftone mask, as we now show.

There are several different morphological ways to identify text and halftones. Some involve binary reconstruction to form the masks at some point in the calculation. The images are assumed to be reasonably well deskewed. Here is an almost trivial approach: do a horizontal closing followed by a smaller horizontal opening. This can leave pixels within text lines as solid fg rectangles, separated vertically by bg pixels,

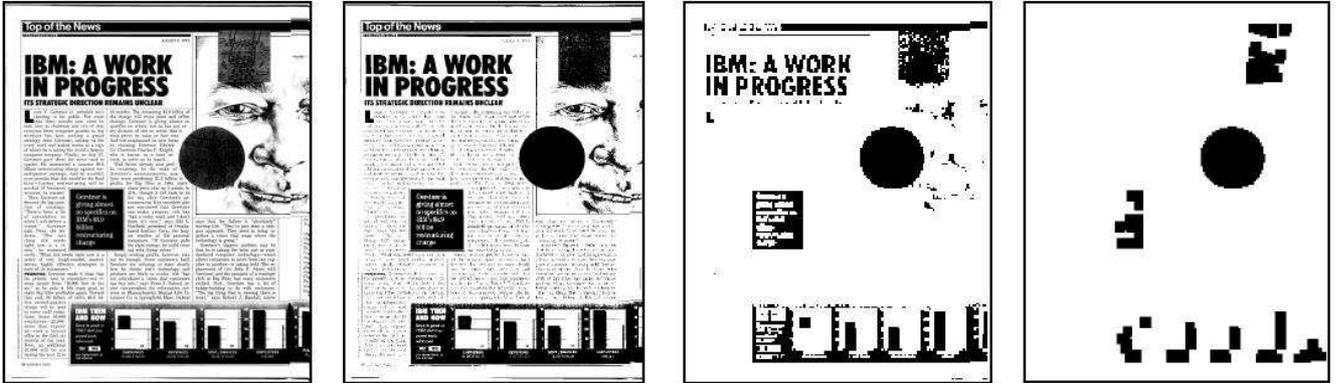


Figure 1: Generation of halftone seed to identify the existence of images.

and pixels within halftone regions as solid fg. This is the essence of an early morphological approach called RLSA [5]. A vertical opening can then remove the text lines, leaving the halftone mask.

We now show a somewhat more accurate method for page segmentation. All operations except the halftone seed construction are performed at a resolution of 150 ppi. Start by finding the binary masks that label image pixels. In the following, we show the operations on two different images that have text, image and rules in nontrivial layouts. Figure 2 shows steps in projecting out the halftone parts of the page (a). The seed (b), composed of pixels that are nearly certain to be within the halftone region(s), is generated by a sequence of 2x rank reductions (levels 4, 4 and 3), followed by a 5×5 opening and 8x replicated expansion back to 150 ppi. This was shown in Figure 1. The clipping mask (c) is designed to connect pixels in each halftone region (so that even a single seed pixel will fill it entirely), but not to form a bridge to any pixels in non-halftone regions. It is generated from (a) using a 2x reduction (level 1) followed by a 4×4 closing. The halftone mask (d) is then generated by binary reconstruction from the seed into the mask.

The next step is to find the text lines. These can be consolidated through a horizontal closing, but such an operation will join lines in different columns, so a vertical whitespace mask must be generated that can later restore the white gutters. This is shown in Figure 3, where in (a) the halftone mask has been subtracted from the original. To build the mask, invert the image (b). Opening with a large vertical Sel can leave components that will break text lines with a large amount of white space above or below, but this can be prevented by opening first with a Sel that is wider than the column separations and higher than the maximum distance between text lines (c). After these pixels are removed, open with a 5×1 horizontal Sel to remove thin vertical lines, followed by opening with a 1×200 vertical Sel to extract long vertical lines (d).

Figure 4 shows the text line extraction process, with the whitespace mask computed in (b). Starting again with the image (a), solidify the text lines using a 30×1 closing (c). Text in adjacent columns that has been joined is then split by subtracting the vertical whitespace mask, and a small 3×3 noise-removal opening yields the textline mask (d).

Figure 5 shows the steps taken to consolidate the text blocks. The original page is shown in (a). Begin with the textline mask, and join pixels vertically using a 1×8 closing (b). Then, for each cc separately, do a 30×30 closing to form a solid mask. By closing separately, we can use a large Sel without danger of joining separate regions. Follow this with a small 3×3 dilation, to insure coverage of the mask components. At this

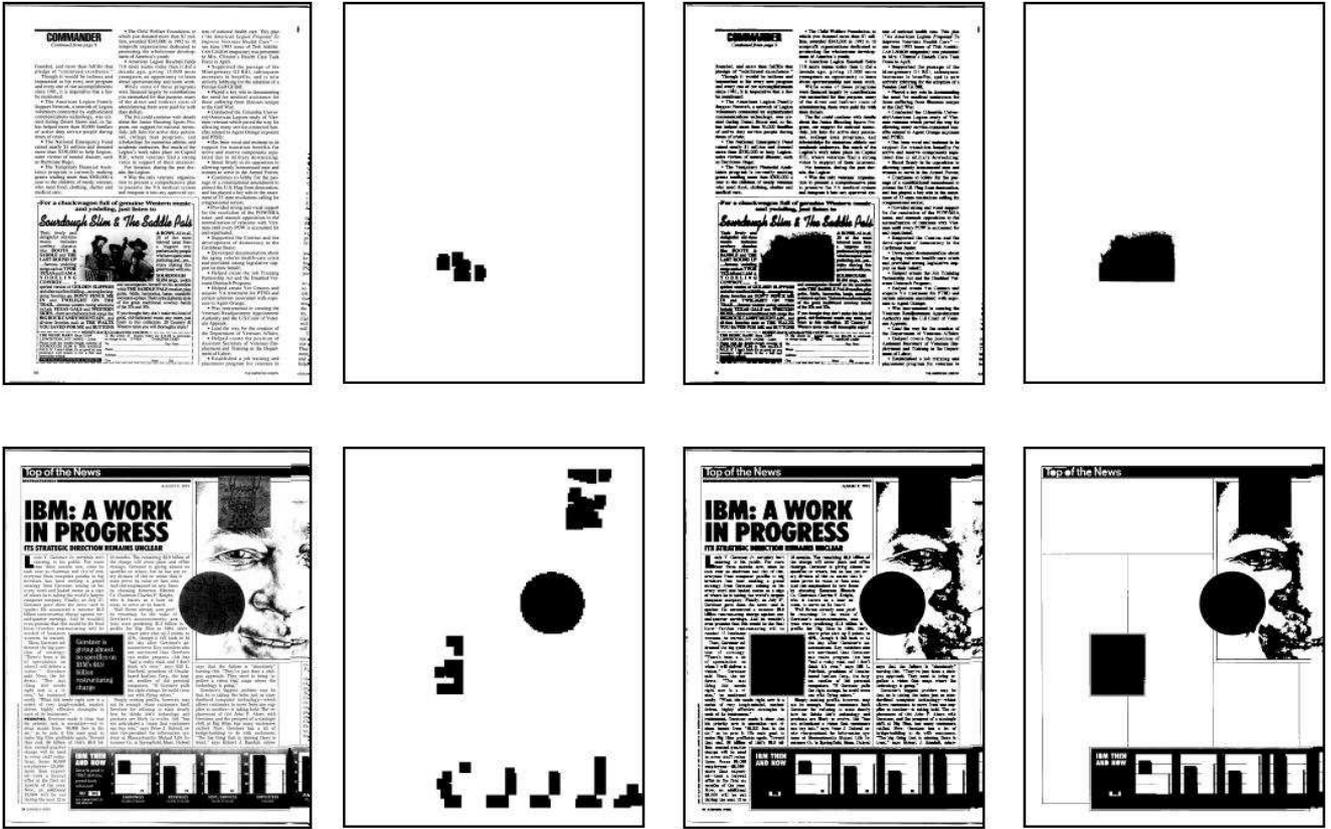


Figure 2: Generation of halftone mask for two different pages.

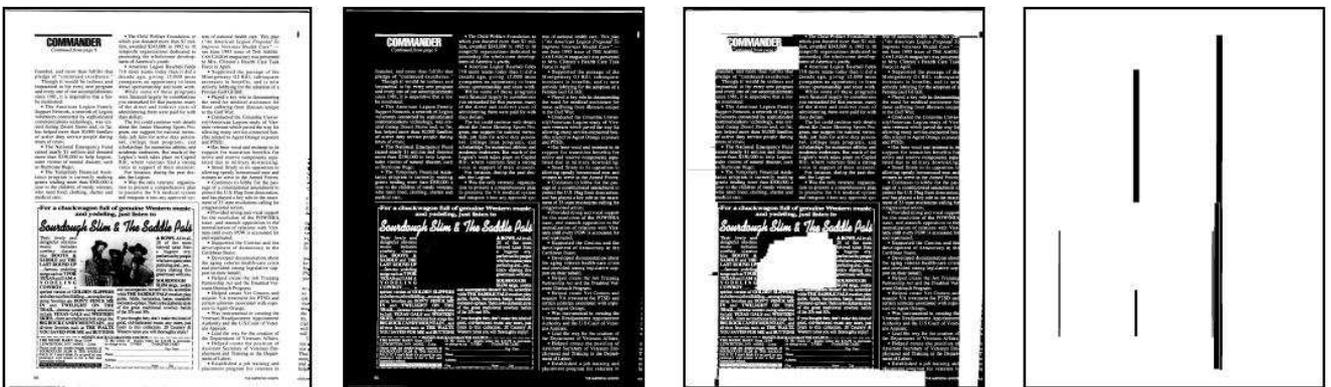


Figure 3: Generation of whitespace mask for example page.

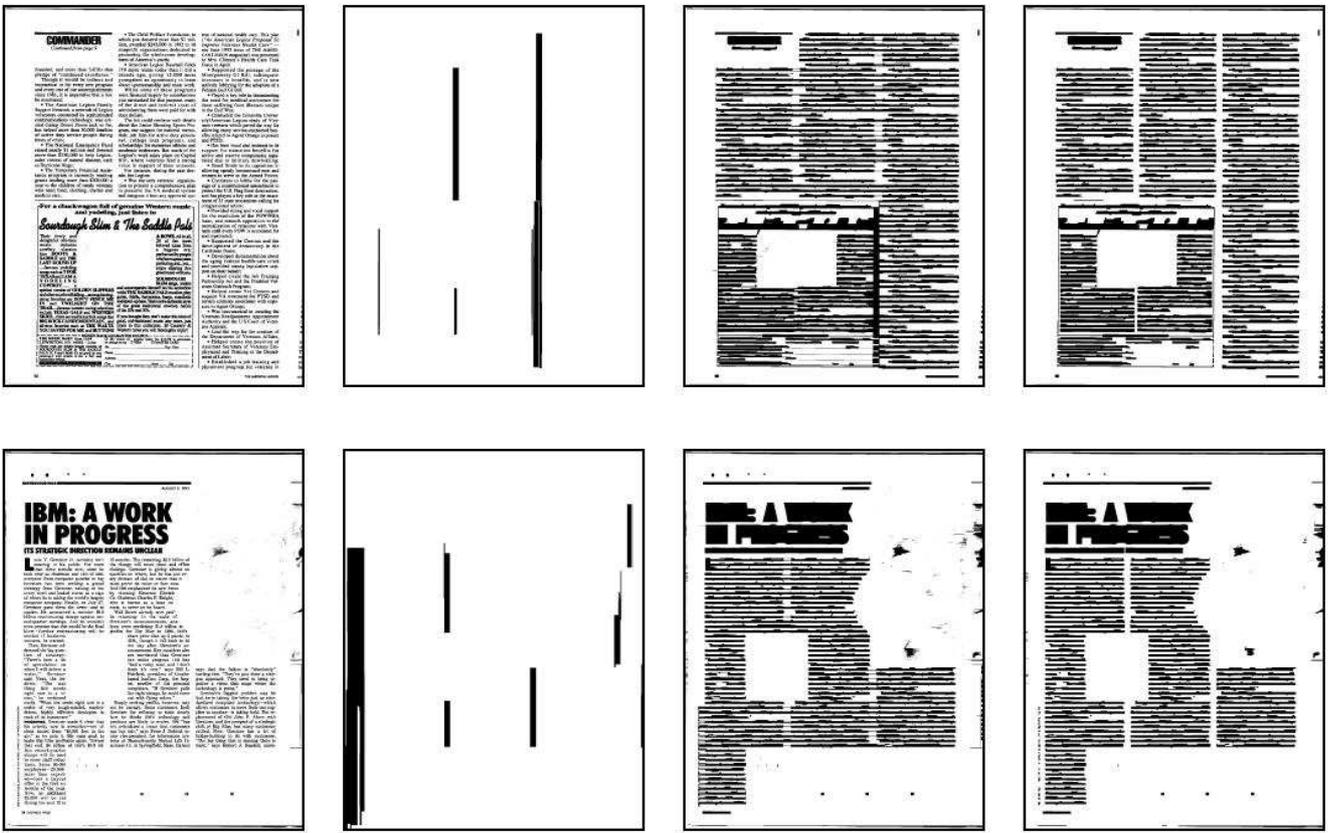


Figure 4: Generation of textline mask for two different pages.

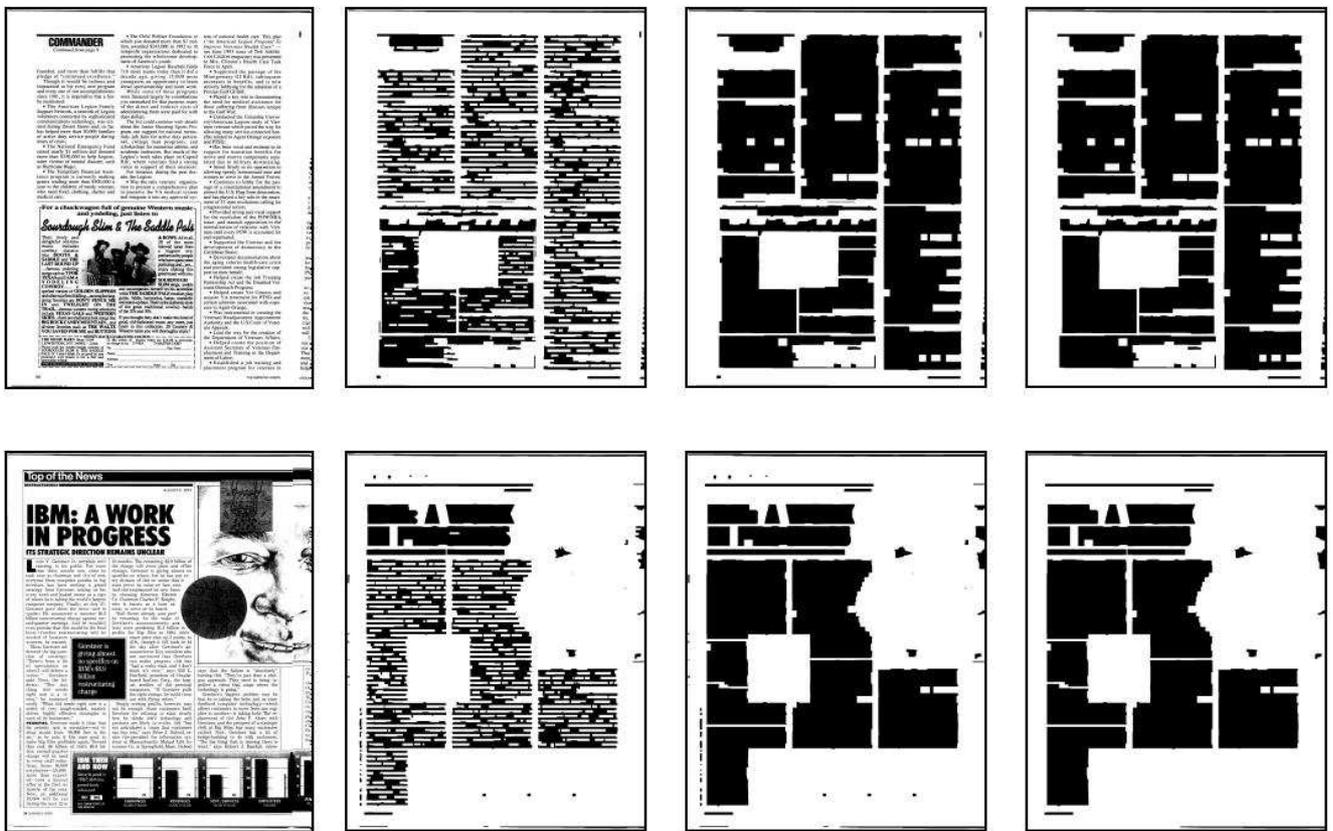


Figure 5: Generation of textblock mask for two different pages.

stage, some textblock components need to be joined horizontally, and this is done with small horizontal closing (c). Because this closing can join textblocks separated by very narrow gutters (which did not happen in the two examples shown), the vertical gutter mask is again applied to split blocks that may have been joined, and small components are removed to obtain the textblock mask (d). This can be further filtered for size and shape.

In these examples of page segmentation, a number of parameters were specified *a priori* for the filter sizes, rather than being computed using measurements on each page. The question naturally arises whether such an open-loop approach is robust. Perhaps surprisingly, the answer is in the affirmative, if by robust we mean that errors where large numbers of pixels are misclassified occur very rarely. The robustness is tested in two ways: (1) by using the algorithm on a large number of pages and (2) by demonstrating that the results are relatively invariant when the parameters are changed by about 30 percent in each direction. The latter is easily measured by scaling the image up and down by this fraction. In this way, it is seen that when computing textblocks on a scaled up image, some of the textlines are not joined, so the vertical closing parameter should be larger. The advantage of this highly-empirical approach is that failures are easy to find and to analyze, and proposed improvements are quickly tested.

2.2 Skew detection

Image deskew greatly simplifies page analysis and improves both the performance of symbol-based compression (*jbig2*) and the displayed appearance of the page. There have been many approaches to skew detection for 1 bpp images, most of which use some variation of a Hough transform or of pixel projection profiles. Others have used fourier transforms, the location of connected components, and special pre-filterings, such as a rosette of morphological pixel correlation filters [13]. For a short description of some of these methods, see [4]. Here, we consider direct computation of pixel sums.

Assume the image has a single, global skew angle, and that there are either horizontal rules or lines of text in the image. When the image is deskewed, some scanlines will have many fg pixels and others will have very few. Consequently, a simple method for finding the skew angle is to rotate the image until the variance of fg pixels on a scanline is maximized. We speak of this measurement as a function of rotation angle as the “signal.” An actual rotation is not necessary; one can either do a vertical shear and sum on rasterlines, or sum directly over pixels on skewed lines.

This approach has four major drawbacks. First, the signal from textlines will have a broad maximum, corresponding to the range of angles through which a raster line can traverse the length of a textline while staying within the x-height. This angular width is approximately the ratio of the x-height to the length of the textline. Second, if a significant fraction of the fg pixels are not text, there will be a large amount of background “noise.” Third, if there are multiple, unaligned columns, the signal will often be weak and misleading, depending on the specific average alignment of the textlines between columns. Finally, the method is fragile when the scan includes part of a second page, particularly if there is a weak signal from the primary page and a strong but skewed signal from the secondary page.

The simplest and arguably the most effective way to avoid these problems was described by Postl[14] in 1988. Instead of maximizing the variance of pixels on a scanline, Postal maximized the variance of the *difference of pixels on adjacent scanlines*. Let the sum of pixels in the i^{th} scanline be $p_i(\theta)$, where θ is the angle through which the image is rotated. Then Postl’s signal is

$$S(\theta) = \sum_i (p_i(\theta) - p_{i-1}(\theta))^2 \quad (1)$$

where the sum extends over all scanlines in the image. The image is then deskewed by rotating through the angle θ for which $S(\theta)$ is maximized. This is effective because, when the page is aligned, most of the signal comes from a relatively small fraction of scanlines; namely, those at the base and x-height of the text lines. Halftone pixels contribute little to such a differential signal. Text lines in each of multiple columns will contribute relatively independently to the signal if they are not aligned. And the peak will be very sharp, corresponding to an angular half-width in radians of approximately $1/(\text{textline width in pixels})$. At 300 ppi, with a textline width of 1500 pixels, the half-width of the peak in $S(\theta)$ is about 0.04 degrees. This is more than sufficient for visual appearance, because it is unusual to notice image skew that is less than 0.2 degrees.

An efficient implementation has several characteristics. It computes at a resolution that meets the accuracy requirements, using the angular estimate described above. It generates low-resolution versions using a cascade of 2x rank reductions with low rank (dilation followed by subsampling), to maintain the signal strength by retaining pixels at the lower resolution. It finds the skew angle with a minimum number of variance measurements, typically using a sweep of angles with equal intervals to locate the peak within about 1 degree, followed by a binary search with 4 or 5 interval halvings. Results have been given on a data set of about 1000 images [2], and these have been compared with a morphologically-based filtering approach [13].

Along with angle corresponding to the maximum score, it is necessary in practice to compute a confidence factor. A reasonable measure of confidence is derived from the ratio of max to min score in the binary search region, along with a threshold on the min score after it is normalized for page size using the product hw^2 .

Suppose the skew is not uniform on the page. This can occur when the scan feeder causes the page to rotate slightly as it is scanned. Then the skew varies approximately linearly with vertical position, and a projective transform is required to remove the skew. The local skew can be found by a set of skew measurements on overlapping horizontal strips, and then doing a linear least squares fit of skew angle to vertical location of the strip. Consider two lines that are near the top and bottom of the page and have the LLS-fitted local skew. These intersect the page sides in four points, which can be used in a projective transform to remove the local skew everywhere.

2.3 Text orientation detection

The *hit-miss transform* (HMT) can be used to determine the orientation of Roman text, because there is a preponderance of ascenders over descenders (approximately 3:1 for English). Consider the four hit-miss Sels:

The hits are black squares, misses are black squares with white circles, don't-cares are white squares, and the origin has a small black circle. The signal in this case is the difference between the number of ascenders, identified from the HMT using the first two Sels, and the number of descenders, using the last two Sels. The statistical significance of this difference is determined as follows. The expected variance in each of these numbers is proportional to their square root. The probability that the two populations can be distinguished (i.e., that the distributions do not overlap) is estimated from the square root of the sum of the individual variances:

$$\sigma_o = \sqrt{N_{up} + N_{down}}/2 \quad (2)$$

Then the *normalized orientation signal* is defined as the difference between the number of ascender and descenders, expressed as a multiple of σ_o :

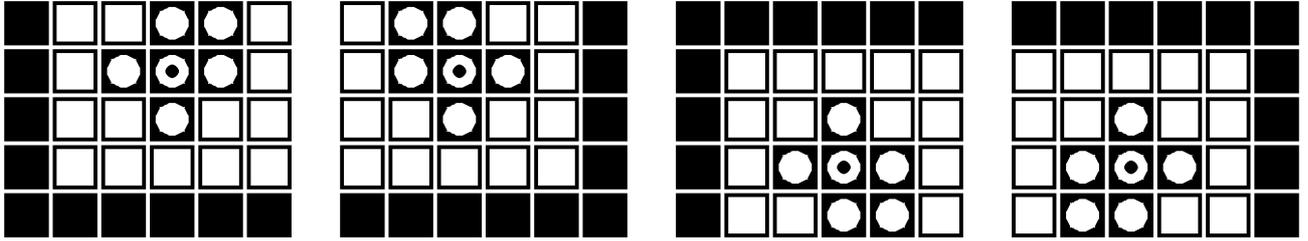


Figure 6: Hit-miss Sels for extracting character ascenders and descenders.

$$\bar{S}_{orient} \equiv |N_{up} - N_{down}| / \sigma_o = 2 |N_{up} - N_{down}| / \sqrt{N_{up} + N_{down}} \quad (3)$$

Usually there will be different prior probabilities for the text orientation, so different thresholds are in general set on the normalized signal for a decision to be made. The signal can also be measured in landscape orientation, and the two signals compared, using appropriate priors, to determine the orientation as one of a set of four directions.

Before doing the HMT, the textline structure should be simplified to fill the holes within the x-height region, leaving only the ascenders and descenders. This can be done with a horizontal closing to solidify the text line, followed by a larger opening to remove all ascenders and descenders that have possibly been joined by the closing. The ascenders and descenders can then be simply reconstructed by ORing with the original image. These pre-HMT operations can usually be done at a lower resolution of between 100 and 150 ppi, using a dilating rank reduction to preserve pixels.

After the HMT we have pixels in small clumps associated with each ascender and descender. To get the ascender and descender count, we can find the number of 8-cc, but a far more efficient and robust way is to do a rank reduction cascade that consolidates each small cluster into a tiny cc (using rank level of 1), followed by counting the number of components at this reduced resolution.

2.4 Word segmentation

The identification of words is useful for many applications. Words are the fundamental unit for generating a reverse index of a document, enabling very rapid search by query. Word images are converted to searchable strings by OCR, but some applications use the word images directly. For example, *document image summarization* (DIMSUM)[6], a very fast extraction of the key words, key phrases, and salient sentences, is enabled by identifying the words. It then performs unsupervised classification on their shapes, analyzes the frequency of words, bigrams and trigrams, and their populations within sentences – all without OCR. Morphological characteristics of the words can be used to identify languages, based on the shape of the most common words, again with little or no OCR required.

The generation of textline and textblock masks in the previous section was a bottom-up merging, starting with the pixels. To find the word bounding boxes, it is simplest to start with a textline and merge the pixels or connected components. This is tricky because the amount of space between words can vary significantly for variable character width fonts, depending on font size and the typesetting algorithm used for right justification. To do a proper image-based segmentation of words, the text lines are sorted by font size, and lines with a similar size are analyzed together.

A simple method for splitting words of roughly comparable font size is to compute the number of cc after each successive dilation with a horizontal 2×1 Sel. The number of cc will quickly fall as the characters within each word are merged, then remain fairly constant as the space between words is reduced, and finally fall again as the words begin to merge. At each iteration, the difference between the number of cc and the number at the previous iteration is found. The iteration number that minimizes this difference gives the optimal dilation, from which the word bounding boxes are derived. For efficiency, this can typically be done at a resolution of about 150 ppi. The method is robust if (1) only text lines of comparable font size are used and (2) the text lines are individually extracted so that there is no possibility of merging text from different text lines.

Figure 7 shows a typical output on a page of text, analyzed at a resolution of 300 ppi to show the details of the distribution. The characters are being joined in the rapid drop through a dilation of 5, and the minimum difference occurs at 7.

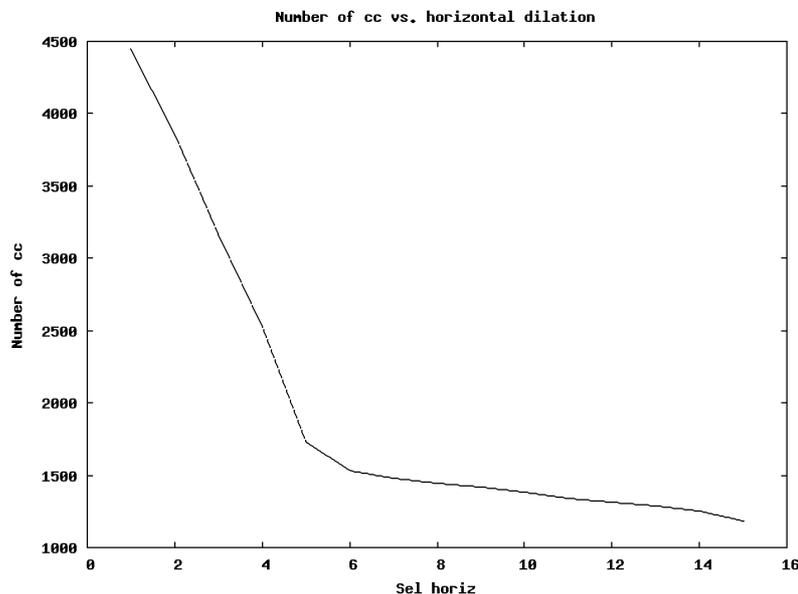


Figure 7: Number of connected components at successive dilations.

2.5 Pattern matching

The ability to do fast pattern matching between elements of document images, such as cc or character or word images, is an important underpinning of many important applications. Some examples are:

- Most OCR systems use image matching with a large library of templates.
- Lossy jbig2 compression of binary images requires unsupervised classification of components into a relatively small number of similarity classes, the templates of which are used to represent each instance of its class when rendering the page.
- The generation of similarity classes can be used to improve the quality of a rendered image, by generating grayscale templates from a set of binary instances. These grayscale templates can be

used directly to substitute for the binary instances, or they can be converted to higher resolution binary templates, a process called “super-resolution.”

- Hit-miss Sels can be generated automatically from a pattern on an image, and then used to find all other occurrences of this pattern.
- Applications such as DIMSUM estimate important words, phrases and sentences by the occurrence of repeated word shapes.

Pattern matching requires some way to measure similarity between elements. Two popular similarity measures for binary images are the Hausdorff distance and correlation. Once a measure is chosen, along with a threshold for declaring two patterns sufficiently similar to belong to the same class and a policy (typically “greedy” or “best match”) for terminating the search for a matching template, unsupervised matching can proceed [11],[16]. We next describe these similarity measures, the methods for implementing them efficiently, and some of the engineering issues for building an unsupervised character classifier from them.

2.5.1 Hausdorff image comparator

The Hausdorff distance H is a true metric (that obeys the triangle inequality) for comparing two 1 bpp images [7]. It is defined as the maximum of two directed Hausdorff distances, h , where the directed Hausdorff distance between images A and B is the maximum over all pixels in A of the distance from that pixel to the closest pixel in B . Formally, if we define the distance from a point p in A to the nearest point in the set B to be $d(p, B)$, then the directed Hausdorff distance from A to B is

$$h(A, B) = \max_{(p \in A)} d(p, B) \quad (4)$$

and the Hausdorff distance is

$$H(A, B) = \max(h(A, B), h(B, A)) \quad (5)$$

The Hausdorff distance is an appealing metric to use in comparing two instances of the same character because we expect most of the pixel variation to occur at the boundary, where the contribution to the distance is small. However, because Hausdorff is sensitive to salt and pepper noise in pixels far from the nearest boundary pixel, it is necessary to use a rank version, with a rank fraction slightly less than 1.0 to give some immunity to such noise [3].

For the classifier application, we have a set of templates for existing classes and a set of instances yet to be assigned to a class (or, if not assigned, to become the template for a new class). Greedy matching works well: each instance must be matched against the templates until a sufficiently close match is found. Instead of computing the Hausdorff distance between two patterns, which is expensive, a decision is simply made whether the distance is less than some threshold, with the rank factor permitting a small number of outliers. The comparison is made for a single alignment, where the patterns have coincident centroids. Then an efficient implementation dilates both patterns in advance, and we check if the dilated image of one contains a rank fraction of pixels in the undilated image of the other, and v.v.

In practice, for small text that is scanned at 300 ppi, character confusion can occur with a Hausdorff distance threshold of 1, which is implemented with dilation by a 3×3 Sel. Consequently, it is necessary to use a 2×2 Sel with a rank fraction of about 0.97. A fraction 0.95 or less results in different characters being placed in the same class; above 0.99 gives too many classes for good compression.

2.5.2 Correlation image comparator

Because very tiny Hausdorff distance thresholds are required to correctly classify small text components, the pixels near the boundary are important. Consequently, correlation comparators, which give equal weight to all pixels and can be more finely tuned, are preferable to rank Hausdorff. The centroids are again aligned when doing the correlation.

Let A and B be the binary images to be compared, and denote the number of fg pixels in an image X by $|X|$ and the number in the intersection of the two images by $|A \oplus B|$. A is one of the templates and B is an instance to be classified. Then the correlation is defined to be the ratio

$$C(A, B) = (|A \oplus B|)^2 / (|A| \times |B|) \quad (6)$$

The correlation is compared with an input threshold. However, because two different thick characters can differ in a relatively small number of pixels, the threshold itself must depend on the fractional fg occupancy of image B . Let the bounding box of B be $w_B \times h_B$. Then the fg occupancy of B is $R = |B| / (w_B \times h_B)$. The modified threshold T' then depends on two input parameters, an input threshold T and a weighting parameter F ($0.0 \leq F < 1.0$):

$$T' = T + (1.0 - T) \times R \times F \quad (7)$$

For 300 ppi images, it is found experimentally that values of $T = 0.8$ and $F = 0.6$ form a reasonable compromise between classification accuracy and number of classes.

2.5.3 Component alignment for substitution

A jbig2 encoder must specify, for each instance in the image, the class membership (an index) and the precise location that the template for that class is to be placed by the decoder. Although the matching score (rank Hausdorff or correlation) is found with centroids aligned, in a significant fraction of instances, the best alignment (correlation-wise) differs by one pixel from centroid alignment. This correction is important for appearance of text, because the eye is sensitive to baseline wobble due to a one-pixel vertical error. It is thus necessary to measure the XOR of the two images at the location where the centroids line up, and at the eight adjacent locations. The best location has the minimum number of pixels in the XOR.

2.5.4 Hit-miss comparator

The HMT is a general filter for matching an arbitrary binary pattern to a binary image. There are no constraints on the content of the pattern fg. However, the characteristics of the hit-miss filter must match the expected variation in the pattern, because the HMT doesn't have a rank parameter: every hit and miss must match. For document images, variation can take the form of boundary noise, salt and pepper noise, rotation, scaling, and other image distortions. As a general rule, it is best to put hits and misses far enough from the boundaries to completely avoid boundary noise. One should avoid using more hits or misses than necessary, because it increases both computation time and the likelihood that an instance is missed. If too few hits or misses are used, false matches will be hallucinated. To reduce sensitivity to small skew and scale changes, the aspect ratio of the pattern should ideally be close to 1.

Here are several methods for automatically generating a hit-miss Sel from a pattern:

- *Run centers.* Form a skeleton of both fg and bg, remove all pixels that are within a specified distance of the boundary, and subsample the remaining points either randomly or along the skeleton. A simple approximation to this is to select a set of vertical and horizontal runs, both fg and bg, and choose the centers of these runs when the centers are not too near the boundary `pixGenerateSelWithRuns()`.
- *Random.* Select pixels randomly, up to given fractions of fg pixels (for hits) and bg pixels (for misses). Do not include any pixels that are within a specified distance of a boundary `pixGenerateSelWithRandom()`.
- *Boundary.* Select a fraction of fg and bg pixels that are at specified distances from the boundary. First the fg and bg contours at the specified distances are generated. Then the hits are chosen by subsampling along a traversal of the fg contour, and likewise for the misses. These four parameters allow flexible specification of the hit-miss Sel `pixGenerateSelBoundary()`.

Figure 8 illustrates a hit-miss Sel generated by the boundary method. The pattern (on top) is reduced 8x and the hits and misses are placed at a distance of 1 from the boundary, with hits subsampled every 6th pixel in the fg and misses every 12th in the bg. The HMT is very fast; on a 25M pixel image, reduced 8x to 400K pixels, it takes about 12 msec.



Figure 8: Pattern and hit-miss Sel generated from it at 8x reduction.

Using just the “T” in the pattern makes the HMT more robust to skew and to variations in scale. Figure 9 shows the pattern and the Sel generated at 4x reduction. The HMT on the 4x reduced image (1.6M pixels) takes 0.2 sec.

2.6 Background estimation for grayscale images

We finish with an application showing the use of grayscale morphology. Suppose a document image is captured in grayscale, but with a significant variation in the background illumination across the page, and you wish to render the image in grayscale but reconstructed as it would appear if the illumination were

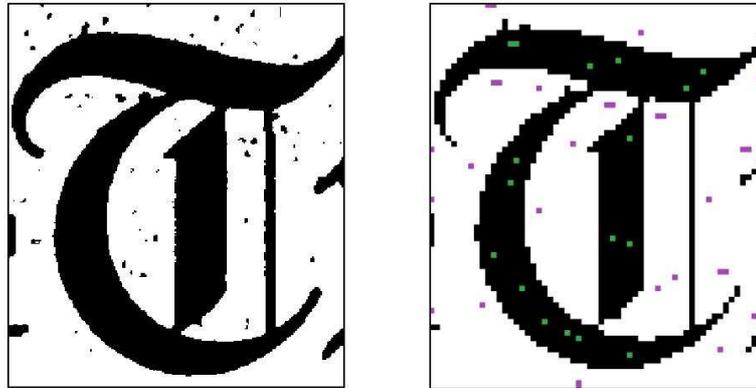


Figure 9: Pattern and hit-miss Sel generated from it at 4x reduction.

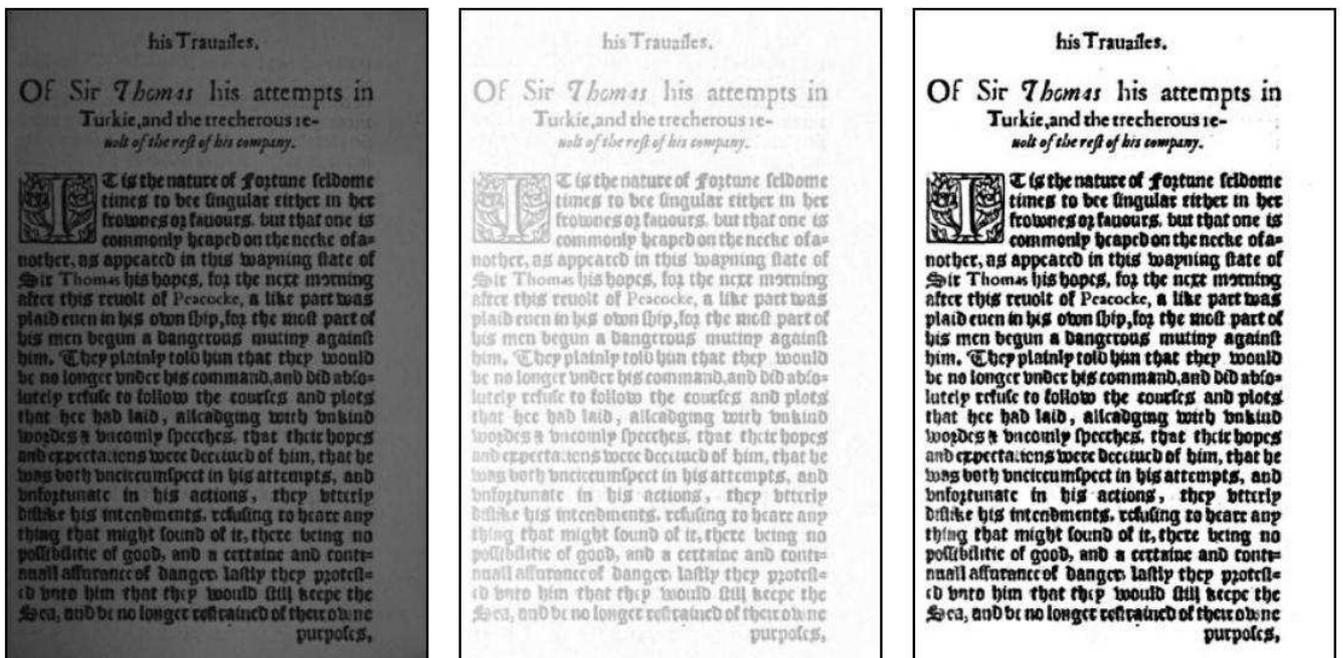


Figure 10: Use of grayscale tophat to compensate for uneven illumination.

uniform. We show two morphologically based approaches that allow more control over the final rendering than simply doing an adaptive threshold to a 1 bpp image.

The first approach uses the morphological tophat directly, where the bg variations are largely removed by first closing the input image (to remove the fg) and then subtracting the input image from the result. Figure 10 shows the processing sequence, starting with an 8 bpp grayscale page image at a resolution of 150 ppi, in (a), and performing a tophat with a 15×15 Sel, which is photometrically inverted (b). The closing in the tophat is performed relatively efficiently using the van-Herk/Gil-Werman (vHGW) algorithm[8, 15], separably, which does the closing in a time independent of the size of the Sel.

The result (b) has a washed-out appearance because the input image (a) has a very dark bg. The appearance can be improved by using a linear *tone reproduction curve* (TRC) to increase the dynamic range, giving (c). In this case, we mapped pixels in (b) with values below 200 to 0 and pixels with values above 245 to 255. The value 245 is chosen for the white point to eliminate most of the bleedthrough from the other side of the page. Nevertheless, the background is not entirely cleaned and the text on the left side of the page is somewhat lighter than the rest.

The second approach also uses the grayscale closing, but it uses the result to control an adaptive grayscale mapping. Figure 11 shows the processing sequence, starting again with the 8 bpp, 150 ppi grayscale page image, in (a). To estimate the background, apply a grayscale closing (max) operation, using a 25×25 Sel. This removes the fg (b), but the blocky residue of the closing is apparent, so we smooth the result using a convolution with a 31×31 flat kernel (c). Like the grayscale closing, the convolution can also be performed in a time independent of the size of the kernel, using an accumulator array of pixel sums over rectangles bounding the upper and left sides of the image. The next step is to multiply the input image (a) by the inverse of image (c). This gives a locally adaptive mapping of the pixels in (a), to compensate for the local illumination. The result (d) should have a fairly uniform background. The appearance can be improved by again applying a linear TRC to increase the dynamic range, giving (e). In this case, we mapped pixels in (d) with values below 30 to 0 in (e), and pixels in (d) with values above 180 to 255 in (e). We can now binarize with a uniform threshold, resulting in the 1 bpp image (f).

Why not simply binarize with an adaptive threshold on (a)? There are two reasons. First, by mapping to a grayscale image, we give ourselves the option to change the gamma and the dynamic range of the image before thresholding. Second, we preserve the option of retaining the mapped grayscale image, which displays better on a screen that supports anti-aliasing.

References

- [1] D. S. Bloomberg, "Image analysis using threshold reduction," *SPIE Conf. 1568, Image Algebra and Morphological Image Processing II*, pp. 38–52, 1991.
- [2] D. S. Bloomberg and G. E. Kopec and L. Dasari, "Measuring document image skew and orientation," *SPIE Conf. 2422, Doc. Rec. II*, pp. 302–316, 1995.
- [3] D. S. Bloomberg and L. Vincent, "Pattern matching using the blur hit-miss transform," *Journal Elect. Imaging*, Vol 9(2), pp. 140–150, April 2000.
- [4] D. S. Bloomberg, "Analysis of document skew," <http://www.leptonica.org/papers/docskew.pdf>.
- [5] K. Wong, R. Casey and F. Wahl, "Document analysis system," *IBM J. Res. Develop*, **26**(2), pp. 647–656, 1982.

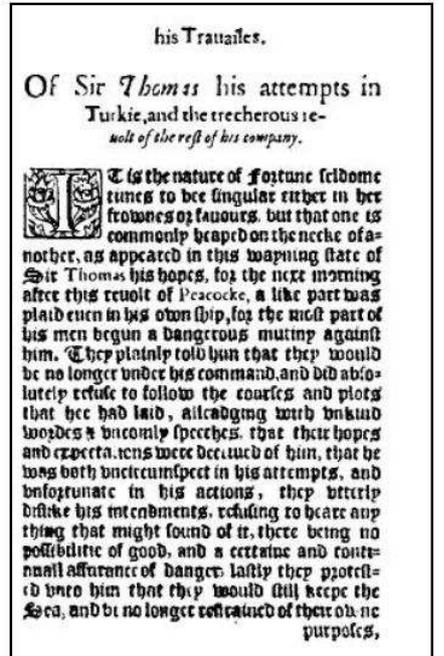
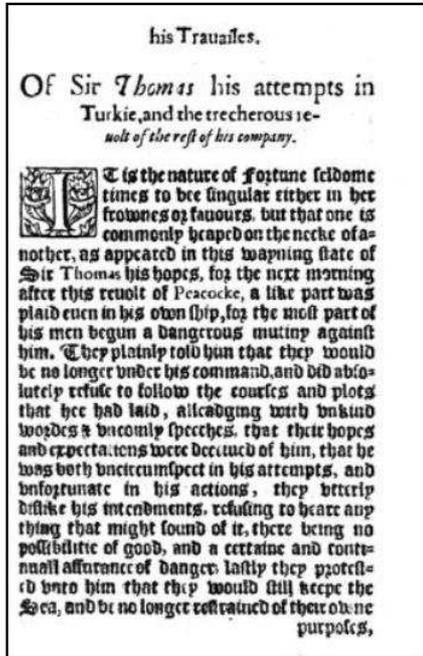
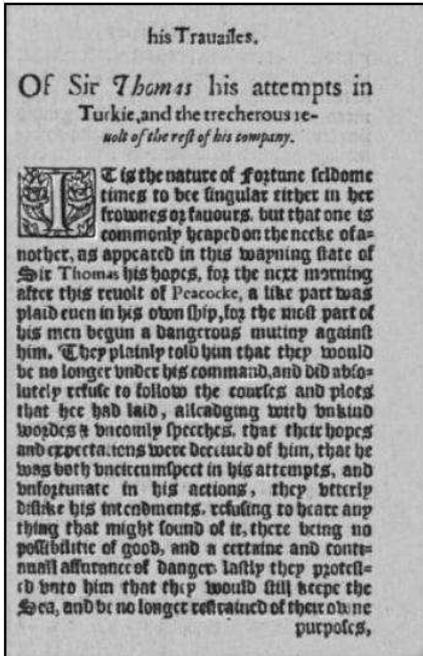
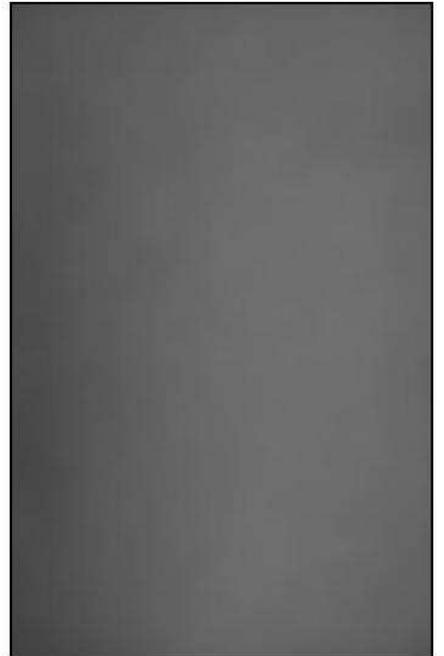
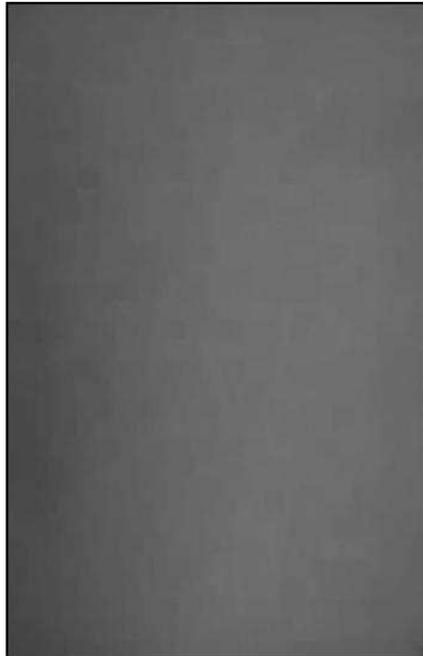
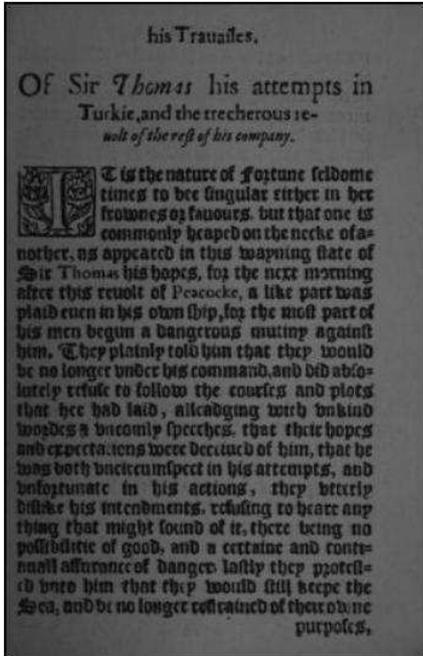


Figure 11: Use of grayscale morphology to estimate bg and compensate for uneven illumination.

- [6] F. R. Chen and D. S. Bloomberg, "Summarization of imaged documents without OCR," *CVIU*, Vol 70, No 3, pp. 307–320, 1998.
- [7] D. Huttenlocher, D. Klanderman, and W. Rucklidge, "Comparing images using the Hausdorff distance," *IEEE Trans. PAMI* **15**, pp. 850–863, Sept. 1993.
- [8] J. Gil and M. Werman, "Computing 2-D min, median and max filters," *IEEE Trans PAMI* **15(5)**, pp. 504-507, May 1993.
- [9] A. Kam and G. Kopec, "Document image decoding by heuristic search," *IEEE Trans. PAMI* **18**, pp. 945–950, Sept. 1996.
- [10] G. Kopec and P. Chou, "Document image decoding using Markov source models," *IEEE Trans. PAMI* **16**, pp. 602–617, June 1994.
- [11] A. G. Langley and D. S. Bloomberg, "Google Books: Making the public domain universally accessible," *SPIE Conf. 6500, Document Recognition and Retrieval XIV*, paper 6500-16, 2007.
- [12] T. P. Minka, D. S. Bloomberg and A. Popat, "Document image decoding using iterated complete path search," *SPIE Conf. 4307, Document Recognition and Retrieval VIII*, pp. 250–258, 2001.
- [13] L. Najman, "Using mathematical morphology for document skew estimation," *SPIE Conf. 5296, Document Recognition and Retrieval XI*, pp. 182–191, 2004.
- [14] W. Postl, "Method for automatic correction of character skew in the acquisition of a text original in the form of digital scan results," *U.S. Pat. 4,723,297*, Feb. 2, 1988.
- [15] M. van Herk, "A fast algorithm for local minimum and maximum filters on rectangular and octagonal kernels," *Patt. Recog. Letters*, **13**, pp. 517-521, 1992.
- [16] www.leptonica.org