Document Marking and Identification using Both Line and Word Shifting *

S. H. Low N. F. Maxemchuk J. T. Brassil L. O'Gorman
AT&T Bell Laboratories
Murray Hill NJ 07974

Abstract

We continue our study of document marking to deter illicit dissemination. An experiment we performed reveals that the distortion on the photocopy of a document is very different in the vertical and horizontal directions. This leads to the strategy that marks a text line both vertically using line shifting and horizontally using word shifting. A line that is marked is always accompanied by two unmarked control lines one above and one below. They are used to measure distortions in the vertical and horizontal directions in order to decide whether line or word shift should be detected. Line shifts are detected using a centroid method that bases its decision on the relative distance of line centroids. Word shifts are detected using a correlation method that treats a profile as a waveform and decides whether it originated from a waveform whose middle block has been shifted left or right. The maximum likelihood detectors for both methods are given.

1 Introduction

Advances in computing, storage and communication technologies have made electronic distribution of publications increasingly popular. The same enabling technologies also tremendously increase the threat of illicit copying and distribution. In [1], a system for secure distribution of electronic documents using a cryptographic protocol is proposed. In [2], we introduce an approach to discourage illicit dissemination by uniquely marking each document copy, so that the original registered recipient can be identified from a recovered illicit copy. The marking – line-shift, work-shift or feature marking – is indiscernible by readers, yet may be reliably identified even from a photocopy of the document. Along the same line but for different medium, schemes are described in [3, 4] to embed unique marking on images to deter illicit copying. Watermarking of optical documents to combat counterfeiting is considered in [5]. General discussion of fingerprinting can be found in [6].

This paper is motivated by an experiment we performed to measure the distortions on a document in the vertical and horizontal directions introduced in printing, photocopying and scanning. The measurements reveal that the distortion is much more severe in one direction than the other. In §2 we model the common distortions and present simple ways to compensate for some of them. In §3 we describe the experiment and offer a plausible explanation of what we observed.

The experiment leads to the marking and detection strategy proposed in §4. A line is marked only if it and its two neighbors are all sufficiently long. The line is marked both vertically using line shifting and horizontally using word shifting. To mark a document by line shift, certain text-lines are shifted slightly, e.g., 1/150 inch, up or down from their normal positions. In word-shift marking, a block of words are shifted slightly to the left or right of their normal position. The shifting pattern is different on different copies. The neighboring lines are not marked. They serve as control for distortion compensation and for estimating the probabilities of detection error on horizontal and vertical profiles. Line shifts or word shifts are detected according as horizontal or vertical profile is less noisy.

In §5 we describe two detection methods. The centroid method measures the distance of the centroid of the marked line to those of its neighboring control lines and decides whether it has been shifted up or down. The correlation method treats a profile as a waveform and decides whether it originated from a waveform whose middle block has been shifted left or right. We motivate the use of centroid detection for line shifts and correlation detection for word shifts by comparing the horizontal and vertical profiles.

*This paper will appear in Infocom'95, Boston, Massachusetts, April 1995. The authors can be reached at {slow, nfm, jtb, log}@research.att.com

We propose a model of detection in [7]. We derive the exact density function of the centroid noise and show that, under certain assumptions generally satisfied by real document profiles, it is approximately zero-mean Gaussian whose variance is computable from the original profile. This allows us to obtain the maximum likelihood detectors for both the centroid and correlation methods and evaluate their relative performance on line and word shifts. We have implemented the strategy proposed in this paper in a software prototype. The prototype is documented and experimental results are presented in [7].

2 Profiles, distortions and compensation

In this section, we model how the profile of an original document is distorted into the profile of its copy by printing, photocopying and scanning. We also present simple methods to compensate for some of the distortions.

The image of a page is represented by a function

$$f(x,y) \;=\; 0 \text{ or } 1, \qquad x \in [0,W], \quad y \in [0,L]$$

where W and L, whose values depend on the scanning resolution, are the width and length of the page, respectively. To simplify notation we assume that x, y and $f(x,y)$ may take continuous values. The image of a text line is simply the function restricted to the region of the text line:

$$f(x,y) \;=\; 0 \text{ or } 1, \qquad x \in [0,W], \quad y \in [t,b]$$

where t and b are the top and bottom 'boundaries' of the text line, respectively. For instance, we may take t or b to be the mid-point of the interline spacing. The *horizontal profile* of the text line

$$h(y) \;=\; \int_0^W f(x,y)\;dx, \qquad y \in [t,b]$$

is the total length of intervals in the x direction on which f is nonzero. The *vertical profile* of the text line

$$v(x) \;=\; \int_t^b f(x,y)\;dy, \qquad x \in [0,W]$$

is the total length of intervals in the y direction on which f is nonzero.

To compensate for certain distortions a line is marked only if it and its two neighboring lines are all sufficiently long. The line can be marked vertically by shifting it slightly up or down from its normal position to carry one bit of the identifier unique to the document. The neighboring lines, called the control lines, are not marked. Alternatively the line can be marked horizontally by shifting certain words slightly left or right from their normal positions. It is divided into three groups of words with a large middle group. The middle group is shifted while its neighbors, called the control groups, remain stationary. Both line-shift and word-shift marking can be considered within the same model where we have a profile that covers three 'blocks'. For line-shift each block is the horizontal profile of a text line. For word-shift each block is the vertical profile of a group of words. The middle block is shifted slightly while the two control blocks are stationary.

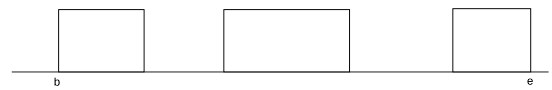

Figure 1: Profile h(x)

Hence consider a three-block profile $h(x)$, $x \in [b,e]$, as shown in Figure 1. Unless specifically noted, $h(x)$ denotes either a horizontal or vertical profile.

Besides 'salt-and-pepper' noise and skewing, the most severe distortions are translation and expansion/shrinkage of the text, blurring, random shifting, and uniform change in intensity. Occasionally from a copier not well maintained, a copy may have a horizontal or vertical strip across the page that is darker or lighter. Such a strip does not significantly affect the centroid of a horizontal profile since each block is typically narrow and tall. A horizontal strip does not affect our word-shift detection since the entire observed horizontal profile $g(x)$ in Proposition 2 in §5.3 below will be multiplied by a constant. A vertical strip does affect word-shift detection and should be compensated for before detection is attempted.

The 'salt-and-pepper' noise and excessive skewing are removed from the bitmap image before the profile is compiled [8, 9]. We model the other distortions as in Figure 2. The profile $h(x)$ is first marked to become $g_0(x)$. The marked but undistorted profile $g_0(x)$ is corrupted by an additive noise into

$$g_1(x) \quad = \quad g_0(x) + N(x), \qquad x \in [b,e]. \qquad (1)$$

Here the random process $N(x)$ models all the noise not accounted for and introduced by distortion compensation to be described. It is translated by l into

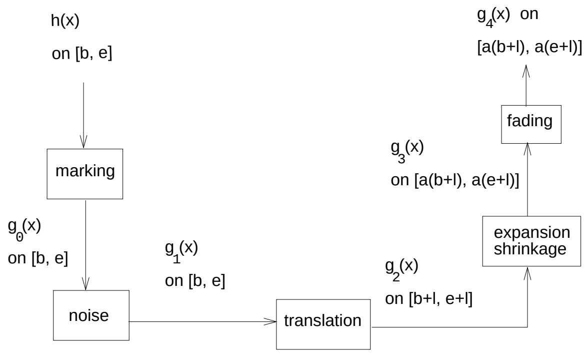

Figure 2: Distortions

the profile

$$g_2(x) \;=\; g_1(x-l), \qquad x \in [b+l, e+l]$$

and then expanded or shrunk by a factor $a > 0$ into

$$g_3(x) \;=\; g_2(\frac{x}{a}), \qquad x \in [a(b+l), a(e+l)].$$

Finally, the profile is faded into the observed profile

$$g_4(x) \;=\; \frac{1}{\gamma(x)} g_3(x), \qquad x \in [a(b+l), a(e+l)]$$

where $\gamma(x) > 0$ is a gain factor.

In the following we describe simple methods to estimate $g_1(x)$ from the observed profile $g_4(x)$. Detecting $g_0(x)$ from the estimate of $g_1(x)$ is discussed in §5.

As noted earlier, $\gamma(x) \neq 1$ only if there is a vertical strip going down the page and $g_4(x)$ under consideration is a vertical profile. In this case, the fading gain can be approximated by a piecewise constant function:

$$\gamma(x) = \sum_i \gamma_i \; 1(x \in [x_i, x_{i+1}))$$

where the indicator function $1(x \in I)$ is 1 if x is in interval I and 0 otherwise. Here, $\{x_i\}$ partitions the interval $[a(b+l), a(e+l)]$. We assume that fading distorts the adjacent control lines in a similar manner. Denote by $c_0(x)$ the profile of the uncorrupted upper control line. It goes through the same distortion process in which $c_0(x)$ is transformed into $c_1(x), ..., c_4(x)$ in a way similar to what $g_0(x)$ in Figure 2 goes through. Then the profile before the fading stage is

$$c_3(x) = c_0(\frac{x}{a} - l) \tag{2}$$

and the observed profile is

$$c_4(x) \;=\; \sum_i \frac{1}{\gamma_i} c_3(x) \; 1(x \in [x_i, x_{i+1})).$$

Integrating both sides over each interval $[x_i, x_{i+1})$, we estimate γ_i by

$$\gamma_i \;=\; \frac{\int_{x_i}^{x_{i+1}} c_3(x)\; dx}{\int_{x_i}^{x_{i+1}} c_4(x)\; dx} \;=\; \frac{aC_0(\frac{x_i}{a} - l, \frac{x_{i+1}}{a} - l)}{C_4(x_i, x_{i+1})} \tag{3}$$

where we have used (2) in the second equality. Here, $C_j(x, y)$, $j = 0, 4$, denotes $\int_x^y c_j(z)dz$. By assumption, the same fading gain γ_i distorts the marked line on each interval $[x_i, x_{i+1})$. Hence for $x \in [a(b+l), a(e+l)]$,

$$g_3(x) \;=\; \sum_i \gamma_i \; g_4(x) \; 1(x \in [x_i, x_{i+1})). \tag{4}$$

The translation l and expansion factor $a > 0$ can be estimated by measuring the beginning position y_0 and the ending position y_1 that defines the three-block profile $g_3(x)$. Then[1]

$$l = \frac{ey_0 - by_1}{y_1 - y_0} \qquad \text{and} \qquad a = \frac{y_1 - y_0}{e - b}. \tag{5}$$

With these, we obtain $g_1(x)$ from $g_3(x)$ by

$$g_1(x) = g_3(a(x+l)), \qquad x \in [b, e]. \tag{6}$$

Denote by I_i the interval $[x_i, x_{i+1})$. Substituting (4) into (6), we estimate the profile $g_1(x)$, $x \in [b, e]$, from the observed profile $g_4(x)$ by

$$g_1(x) \;=\; \sum_i \gamma_i g_4(a(x+l)) \; 1(a(x+l) \in I_i) \tag{7}$$

where a and l are given by (5), and γ_i is given by (3).

3 A noise experiment

In this section, we describe the noise experiment that motivates our marking and identification strategy, present some of the measurements, and give a brief explanation. The implication of the results is discussed in the next section.

We took a page of text consisting of alternating symbols of '|' and '-' on each line, copied ten times, scanned in each copy to produce a bitmap image. The bitmap image was processed to correct for excessive skewing (rotation of text) and to remove 'salt-and-pepper' noise. Then the vertical profile of each textline is compiled. We compensated for the translation, expansion/shrinkage, and fading using the methods

[1] The effect of translation is explicitly accounted for in the centroid detector by using the distance between adjacent centroids for detection (see §5.2 below).

described in §2. We performed the experiment for both photocopies in the portrait and landscape orientation. When photocopied in the portrait orientation, the page consisted of 27 lines, each line having 40 '|'s, giving a sample size of 1080 '|'s; for landscape orientation, it had 22 lines, each line having the same number of '|'s, giving a sample size of 880 '|'s.

Unlike a typical profile for regular English text, the vertical profile of each text-line has easily identifiable columns corresponding to vertical strokes '|'s. We will refer to the laser printed copy as the 0th copy; the ith copy, $i = 1, ..., 10$, was produced by photocopying copy $i-1$. For stroke j on the ith copy, $i \geq 0$, we measured its 'distance' d_{ij} relative to the left text-margin. Here, $j = 1, .., 1080$ for portrait-style photocopies, and $j = 1, .., 880$ for landscape-style photocopies. Since a stroke, or the corresponding column in the profile, is typically 3-5 pixels wide, this 'distance' is taken to be the centroid of the column. Then

$$n_{ij} = d_{ij} - d_{0j}, \qquad i \geq 1$$

measures the displacement of the jth stroke on the ith copy from its original position. For each copy i, we calculated the sample mean and variance by averaging over j. They are plotted in Figure 3 as a function of copy generation. It shows that the mean displacement on the profile copied in portrait orientation is about 2.5 times that in the landscape orientation. More importantly the variance in the portrait orientation is about twice as much as that in the landscape orientation.

That the noise is much more severe in one direction than the other becomes intuitive in retrospect. Major sources of distortion are printing, photocopying, scanning and subsequent processing of the bitmap image and profile. What we observed was due to the remaining noise after 'salt-and-pepper' noise, excessive skewing, translation and expansion/shrinkage of entire text, and fading have been compensated for. Apparently, noise introduced by the copier dominated that from other sources. Careful examination of the profile indicated that random shifting of words within the text is more severe in the portrait than landscape copies, leading to our experimental result. The reason for this lies in the electrophotographic process in the copier. We shall sketch below a plausible reason; see [10] for a detailed discussion on the physics of photocopying and [11] on image defect models.

The eletrophotographic process takes six distinct steps. The first three steps produce an image of the document on a photoreceptor. This image is transferred from the photoreceptor onto paper in the forth

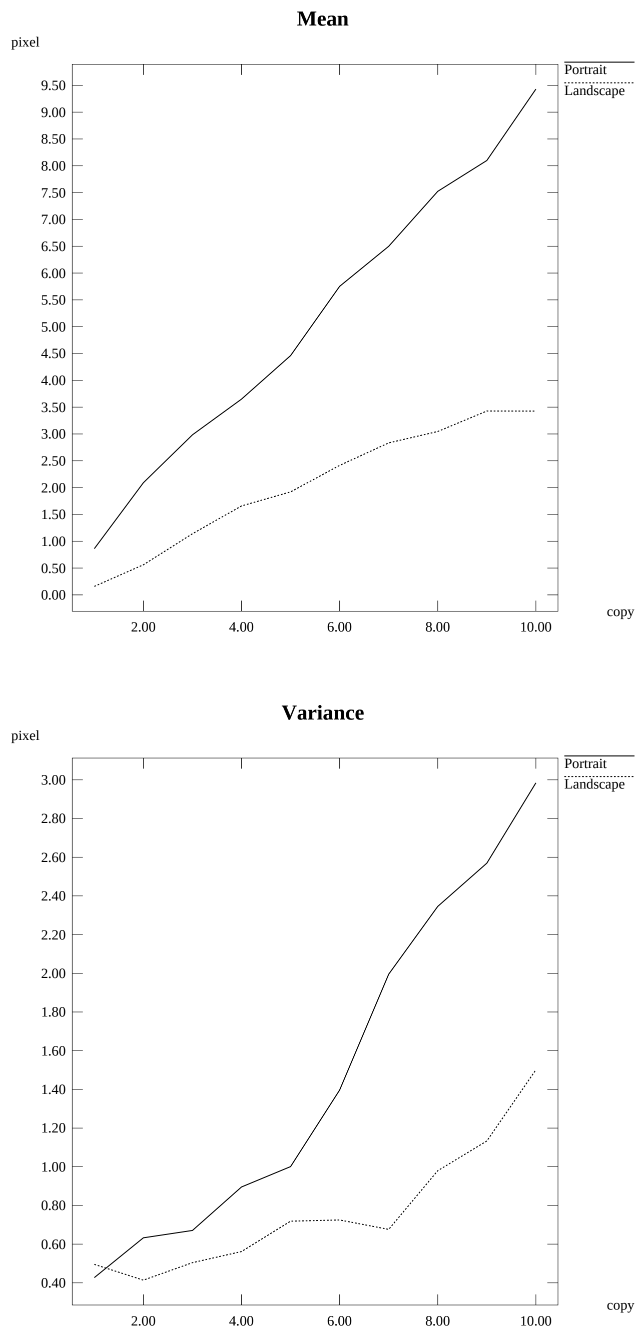

Figure 3: Mean and variance of displacement

and fifth step. In the last step, the photoreceptor is cleaned for the next copying, which repeats these six steps. These six steps are done at different parts of the copier. In most copiers, the photoreceptor material is mounted on a drum or a conveying belt which rolls through the six different points of the copier where these steps are performed. The paper rolls from the supply paper tray through where the forth and fifth steps are performed, and then out of the copier. The random shifting of the image is more severe in the direction of the paper path, which is also the direction in which the photoreceptor moves, than in the orthogonal direction. For our copier and the copying option we selected, this direction happened to be the horizontal direction on the text. This plausibly explains what we observed in the experiment.

4 Marking and detection strategy

In this section we propose a marking and detection strategy which takes advantage of the possibility that the horizontal and vertical profiles can be distorted to different degrees.

Each document copy to be sent to a registered recipient is assigned a unique binary identifier. Certain text-lines that satisfy the first condition below are selected to be marked. Each such line carries one bit of the identifier. When an illicit copy is recovered marking is detected from these text-lines, identifying the registerred recipient.

Marking:

1. A text line is marked only if it and its two neighbors are all sufficiently long. The neighboring lines, called the control lines, are unmarked. They are used for distortion compensation as described in §2 and estimation of the remaining noise.

2. A marked line is marked both vertically using line shifting and horizontally using word shifting. In line shifting, it is shifted slightly up or down from its normal position. In word shifting, it is divided into three blocks of words and the middle block is shifted slightly left or right whereas the two neighboring blocks remain stationary. In either case the *direction* of the shift is determined by the unique identifier.

Detection:

1. Scan in the recovered copy, and the original unmarked copy if its bitmap is not already available.

2. Compile the horizontal and the vertical profile from the bitmap image.

3. Compensate for fading, translation and expansion/shrinkage of the entire text using (7) in §2.

4. Estimate the probabilities of detection error on the horizontal and the vertical profiles.

5. Detect line shift or word shift according as the horizontal or vertical profile is less noisy.

We have tried two ways to estimate the error probabilities in Step 4 above but have not had enough practical experience to tell which is better. The first method marks the control lines on a page both horizontally and vertically with a predefined shifting pattern that is the same for all copies. Every other control line is shifted up or down. These vertical shifts are detected using the centroid detector described in the next section. The success rate is an estimate of the detection error on the horizontal profiles. In the other direction each control line is divided into three groups of words. The middle group is shifted left or right and the neighboring control groups remain stationary. These word shifts are detected using the correlation detector described in the next section. The success rate is an estimate of the detection error on the vertical profiles.

The second method compare the expressions derived in [7] for error probabilities in centroid and correlation detection. These probabilities depend on the profiles of the unmarked original as well as the variance of the additive noise $N(x)$ in (1). The noise variance is estimated from the control lines.

We next describe methods to detect line and word shifts.

5 Detection

In this section we describe two detection methods. The centroid method measures the distance between the centroid of adjacent blocks and decides whether the middle block has been shifted up or down. The second method treats the profile as a waveform and decides whether it originated from a waveform corresponding to a left or right shifted middle block.

We motivate the use of centroid method for line shifts and correlation method for word shifts by comparing the horizontal and vertical profiles. A more formal justification is given in [7].

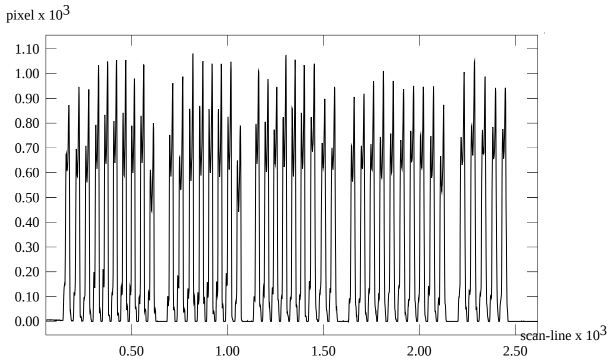

Figure 4: Horizontal profile of a page (resolution = 300 dots-per-inch)

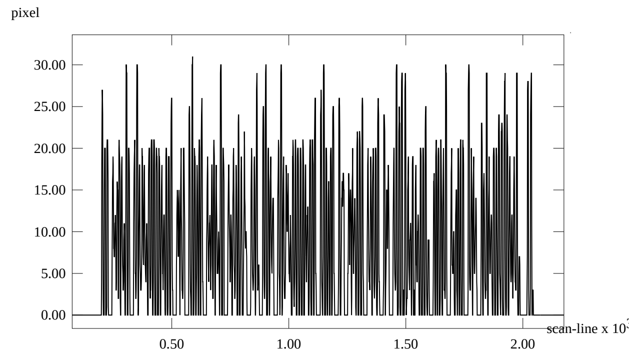

Figure 5: Vertical profile (resolution = 300 dots-per-inch)

5.1 Vertical vs. horizontal profiles

A horizontal profile typically consists of distinct 'columns' and 'valleys'. The 'columns' correspond to text lines and the 'valleys' to interline spaces. Figure 4 shows such a profile of the forth copy generation of a page. The bulk of a column is several hundred pixels for the shown scanning resolution. Distortions introduced in printing, photocopying and scanning typically has a magnitude of 15 pixels or less for this example, an order of magnitude smaller. Hence a simple threshold scheme is sufficient to determine the left and right edges that define a column in the profile. Moreover, since the threshold (on the order of 20 pixels) is small compared with the bulk of a column, an error of a couple scan lines in delineating the column has insignificant effect on the computed centroid. Hence, the centroid method has been very reliable.

One is first tempted to use the same procedure to detect word shifts, with 'columns' in the vertical profile corresponding to words and 'valleys' in between corresponding to interword spaces. The fact that a vertical profile is much shorter and its 'columns' and 'valleys' much less distinguishable makes this approach ineffective. Figure 5 shows the vertical profile of a line on the same photocopy used in generating Figure 4. The 'columns' are almost all less than 30 bits, and the 'valleys' much narrower than in the horizontal profile. It is difficult to accurately identify a 'column' from such a profile. Because of the short profile, an inaccuracy of even a few pixels in delineating such a column (word segmentation) will significantly affect the computed centroid. In addition the centroid noise has a larger variance for a shorter and wider column as shown in [7].

In effect, in line-shift detection, we approximate a horizontal profile by a series of delta functions, one at the centroid of each column, and make our decision based on this approximation. The experimental results in [2] and [7] attest to the accuracy of this approximation for detection purpose. Comparison of Figures 4 and 5 suggests that such an approximation will be poor for a vertical profile.

5.2 Line-shift detection

The centroid method works well only if each block of the given profile can be accurately delineated. Assume this has been done and we have a profile $h(x)$ and three intervals $[b_1, e_1]$, $[b_2, e_2]$, and $[b_3, e_3]$ that define the three blocks in $h(x)$. This is the original unmarked profile. The centroid of block i, $i = 1, 2, 3$, is

$$c_i \;=\; \frac{\int_{b_i}^{e_i} x h(x)\, dx}{\int_{b_i}^{e_i} h(x)\, dx}.$$

We also have the profile

$$g(x) \;=\; h(x) \;+\; N(x), \qquad x \in [b_1,\ e_3] \qquad (8)$$

of the marked copy on the same interval $[b_1, e_3]$. Here, the additive noise $N(x)$ models all the distortions not accounted for and introduced by distortion compensation described in §2. The control blocks have centroids

$$U_1 \;=\; c_1 \;+\; V_1 \qquad \text{and} \qquad U_3 \;=\; c_3 \;+\; V_3.$$

The middle block has been shifted by a size $\epsilon > 0$ so that its centroid is

$$U_2 \;=\; c_2 \;+\; V_2 \;-\; \epsilon$$

if it is up shifted, and

$$U_2 \;=\; c_2 \;+\; V_2 \;+\; \epsilon$$

if it is down shifted. Here, V_i, $i = 1, 2, 3$, are random noises that corrupt the centroids c_i. To eliminate the effect of translation we base our detection on the distance $U_i - U_{i-1}$ between adjacent centroids instead of centroid U_2 of the middle block. Hence we have a classical detection problem in which we have to decide whether the middle block has been up or down shifted given the observed values of $U_2 - U_1$ and $U_3 - U_2$.

Under the assumption that the additive noise $N(x)$ in (8) is white zero-mean Gaussian, we show in [7] that the centroid noise V_i is approximately zero-mean Gaussian whose variance ν_i^2 is a function of the uncorrupted profile $h(x)$. We also derive there the maximum likelihood detector that chooses the direction of the shift that is most likely to have caused the observed $U_2 - U_1$ and $U_3 - U_2$. It minimizes the average probability of decision error when the middle block is equally likely to have been shifted up or down a priori; see, e.g., [12].

It is convenient to use as decision variable the differences

$$\begin{aligned} \Gamma_u &:= (U_2 - U_1) - (c_2 - c_1) \\ \Gamma_l &:= (U_3 - U_2) - (c_3 - c_2) \end{aligned}$$

of the corrupted centroid separations and the uncorrupted separations. Γ_u is the change in the distance of the middle block from the upper control block and Γ_l is that from the lower control block. Without noise $\Gamma_u = -\epsilon$ and $\Gamma_l = \epsilon$ if the middle block is up shifted, and $\Gamma_u = \epsilon$ and $\Gamma_l = -\epsilon$ if it is down shifted. Hence it is reasonable to decide that the middle block is up shifted if $\Gamma_u \leq \Gamma_l$, and down shifted otherwise. With noise, according to the following proposition from [7], these changes in the distance of the middle block from the control blocks should be weighted by the noise variances in the centroids of the control blocks before comparing them. Note that the decision does not depend on the middle block, except through Γ_u and Γ_l.

Proposition 1 *The maximum likelihood detector, when the observed value of (Γ_u, Γ_l) is (γ_u, γ_l), is*

$$\begin{array}{ll} \textit{decide up shift} & \textit{if } \frac{\gamma_u}{\nu_1^2} \leq \frac{\gamma_l}{\nu_3^2} \\ \textit{decide down shift} & \textit{otherwise} \end{array}$$

where ν_1^2 and ν_3^2 are the centroid noise variances of the upper and lower control blocks, respectively.

5.3 Word-shift detection

We again have a profile $h(x)$ and three intervals $[b_1, e_1]$, $[b_2, e_2]$, and $[b_3, e_3]$ that define three groups of words. We assume that $h(x) = 0$ between these intervals. The middle group is shifted slightly while controls groups remain stationary. Let $h^l(x)$ be the resultant profile when the middle block is left shifted by $\epsilon > 0$:

$$h^l(x) = \begin{cases} h(x), & x < b_2 - \epsilon \text{ or } x > e_2 \\ h(x+s), & b_2 - \epsilon \leq x \leq e_2 - \epsilon \\ 0, & e_2 - \epsilon \leq x < e_2 \end{cases} \quad (9)$$

and $h^r(x)$ be that when the middle block is right shifted:

$$h^r(x) = \begin{cases} h(x), & x < b_2 \text{ or } x > e_2 + \epsilon \\ 0, & b_2 \leq x < b_2 + \epsilon \\ h(x-s), & b_2 + \epsilon \leq x \leq e_2 + \epsilon \end{cases} \quad (10)$$

Naturally we assume the shift s is smaller than the interblock spacing. As explained in §2 the profile $g(x)$ compiled from the illicit copy and after distortion compensation is corrupted by an additive noise $N(x)$ such that

$$g(x) = h^l(x) + N(x), \qquad y = b_1, \ldots, e_3$$

if the middle block is left shifted, and

$$g(x) = h^r(x) + N(x), \qquad y = b_1, \ldots, e_3$$

if it is right shifted.

We have to decide whether the middle block is left or right shifted based on the observed profile $g(x)$. Under the assumption that $N(y)$ is zero-mean white Gaussian noise, standard procedure [12, Chapter 4] leads to the following result.

Proposition 2 *The maximum likelihood detector given the observed profile $g(y)$ is:*

$$\begin{array}{ll} \textit{decide left shift} & \textit{if } \sum_{b_1}^{e_3} g(y)\,(h^l(y) - h^r(y)) \geq 0 \\ \textit{decide right shift} & \textit{otherwise} \end{array}$$

where h^l and h^r are computed from the original profile $h(x)$ according to (9) and (10), respectively.

6 Conclusion

We have presented an experiment which showed that, depending on the processes a document goes through, it could be distorted in the vertical and horizontal directions to very different degrees. To take advantage of such possibility we have proposed a strategy that marks a text line both vertically using line shifting and horizontally using word shifting. Probabilities of detection error on horizontal and vertical profiles are estimated from the control lines accompanying each marked line. Line shifts or word shifts are detected according as horizontal or vertical profiles are less noisy. The maximum likelihood detectors have been given for the centroid and the correlation methods.

This marking and detection strategy has been implemented in a software prototype [7]. Preliminary experimental results are encouraging. With a size of 2 pixels (1/150 inch) for both line and word shifts, the system has correctly detected all line and word shifts for photocopies up to the 10th generation in the landscape orientation. For portrait copies, all line shifts has been correctly detected for up to the 10th generation but word-shift errors occurred after the 2nd generation. See [7] for details.

References

[1] A. K. Choudhury, N. F. Maxemchuk, S. Paul, and H. G. Schulzrinne. Copyright protection for electronic publishing over computer networks. Technical Memo BL011382-940428-75TM, AT&T Bell Laboratories, April 1994. Submitted for publication.

[2] J. Brassil, S. Low, N. Maxemchuk, and L. O'Gorman. Electronic marking and identification techniques to discourage document copying. *Proceedings of Infocom'94*, pages 1278–1287, June 1994.

[3] K. Tanaka, Y. Nakamura, and K. Matsui. Embedding secret information into a dithered multilevel image. *Proceedings of the 1990 IEEE Military Communications Conference*, pages 216–220, September 1990.

[4] Germano Caronni. Assuring ownership rights for digital images. Submitted to ASIACRYPT'94, 1994.

[5] Rudolf L. van Renesse, editor. *Optical Document Security*. Artech House, 1994.

[6] Neal R. Wagner. Fingerprinting. *Proceedings of the 1983 IEEE Symposium on Security and Privacy*, pages 18–22, April 1983.

[7] S. H. Low, A. M. Lepone, and N. F. Maxemchuk. Document marking and identification: a model and an implementation. Submitted for publication, December 1994.

[8] L. O'Gorman. Image and document processing techniques for the RightPages Electronic Library System. *Int. Conf. on Pattern Recognition (ICPR)*, pages 260–263, September 1992.

[9] L. O'Gorman. The document spectrum for structural page layout analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11), November 1993.

[10] L. B. Schein. *Electrophotography and Development Physics*. Springer-Verlag, second edition, 1992.

[11] H. S. Baird. Document image defect models. In H. S. Baird, H. Bunke, and K. Yamamoto, editors, *Structured Document Image Analysis*, pages 546–556. Springer-Verlag, Berlin, 1992.

[12] Harry Van Trees. *Detection, Estimation, and Modulation Theory*, volume I. John Wiley & Sons, 1968.