Data Hiding in Document Images Using Digital Watermarking

1. Prof.P.B.Khatkale, 2.Prof. K.P.Jadhav, 3. Prof. M.V.Khasne
1. Lect in K.B.P.Polytechnic, Kopargaon
2. Lect in K.B.P.Polytechnic, Kopargaon
3. Lect in K.B.P.Polytechnic, Kopargaon

Abstract

With the proliferation of digital media such as images, audio, and video, robust digital watermarking and data hiding techniques are needed for copyright protection, copy control, annotation, and authentication of document images. While many techniques have been proposed for digital color and gray scale images, not all of them can be directly applied to binary images in general and document images in particular. The difficulty lies in the fact that changing pixel values in a binary image could introduce irregularities that are very visually noticeable. Over the last few years, we have seen a growing but limited number of papers proposing new techniques and ideas for binary image watermarking and data hiding.

1. Introduction

Given the increasing availability of cheap yet high quality scanners, digital cameras, digital copiers, printers and mass storage media the use of document images in practical applications is becoming more widespread. However, the same technology that allows for creation, storage and processing of documents in digital form, also provides means for mass copying and tampering of documents. Given the fact that digital documents need to be exchanged in printed format for many practical applications, any security mechanism for protecting digital documents would have to be compatible with the paper-based infrastructure. Consider for example the problem of authentication. Clearly an authentication tag embedded in the document should survive the printing process. That means that the authentication tag should be embedded inside the document data rather than appended to the bit stream representing the document. The reason is that if the authentication tag is appended to the bit stream, a forger could easily scan the document, remove the tag, and make changes to the scanned copy and then print the modified document. The process of embedding information into digital content without causing perceptual degradation is called data hiding. A special case of data hiding is digital watermarking where the embedded signal can depend on a secret key. One main difference between data hiding and watermarking is in whether an active adversary is present. In watermarking applications like copyright protection and authentication, there is an active adversary that would attempt to remove, invalidate or forge watermarks. In data hiding there is no such active adversary as there is no value associated with the act of removing the hidden information. Nevertheless, data hiding techniques need to be robust against accidental distortions. A special case of data hiding is steganography (meaning covered writing in Greek), which is the science and art of secret communication. The most important issue in steganography is that the very presence of a hidden message must be concealed. Such a requirement is not critical in general data hiding and watermarking problems.

2. Applications of Data hiding:

Before we describe the different techniques that have been devised for data hiding, digital watermarking and steganography for document images, we briefly list different applications that would be enabled by such techniques.

1. Ownership assertion: To assert ownership of a document, Alice can generate a watermarking signal using a secret private key, and embed it into the original document. She can then make the watermarked document publicly available. Later, when Bob contends the ownership of a copy derived from Alice's original, Alice can produce the unmarked original and also demonstrate the presence of her watermark in Bob's copy. Since Alice's original is unavailable to Bob, he cannot do the same provided Alice has embedded her watermark in the proper manner [14]. For such a scheme to work, the watermark has to survive operations aimed at malicious removal. In addition, the watermark should be inserted in such a manner that it cannot be forged, as Alice would not want to be held accountable for a document that she does not own [10].

2. Fingerprinting: In applications where documents are to be electronically distributed over a network, the document owner would like to discourage unauthorized duplication and distribution by embedding a distinct watermark (or a fingerprint) in each copy of the data. If, at a later point in time, unauthorized copies of the document are found, then the origin of the copy can be determined by retrieving the fingerprint. In this application the watermark needs to be invisible and must also be invulnerable to deliberate attempts to forge, remove or invalidate. The watermark should also be resistant to collusion. That is, a group of k users with the same document but containing different fingerprints should not be able to collude and invalidate any fingerprint or create a copy without any fingerprint.

3. Copy prevention or control: Watermarks can also be used for copy prevention and control. For example, every copy machine in an organization can include special software that looks for a watermark in documents that are copied. On finding a watermark the copier can refuse to create a copy of the document. In fact it is rumored that many modern currencies contain digital watermarks which when detected by a compliant copier will disallow copying of the currency. The watermark can also be used to control the number of copy generations permitted. For example a copier can insert a watermark in every copy it makes and then it would not allow further copying when presented a document that already contains a watermark.

4. Authentication: Given the increasing availability of cheap yet high quality scanners, digital cameras, digital copiers and printers, the authenticity of documents has become difficult to ascertain. Especially troubling is the threat that is posed to conventional and well established document based mechanisms for identity authentication, like passports, birth certificates, immigration papers, driver's license and picture IDs. It is becoming increasingly easier for individuals or groups that engage in criminal or terrorist activities to forge documents using off-the-shelf equipment and limited resources. Hence it is important to ensure that a given document 234 [17]. Was originated from a specific source and that it has not been changed, manipulated or falsified. This can be achieved by embedding a watermark in the document. Subsequently, when the document is checked, the watermark is extracted using a unique key associated with the source, and the integrity of the data is verified through the integrity of the extracted watermark. The watermark can also include information from the original document that can aid in undoing any modification and recovering the original.

5. Metadata Binding: Metadata information embedded in an image can serve many purposes. For example, a business can embed the Web site URL for a specific product in a picture that shows an advertisement for that product. The user holds the magazine photo in front of a low-cost CMOS camera that is integrated into a personal computer, cellular phone, or a personal digital assistant. The data are extracted from the low-quality picture and is used to take the browser to the designated Web site. For example, in the media bridge application (http://www.digimarc.com), the information embedded in the document image needs to be extracted despite distortions incurred in the print and scan process. However, these distortions are just a part of a process and not caused by an active and malicious adversary. Over the last few years, a variety of digital watermarking and data hiding techniques have been proposed for such purposes. However, most of the methods developed today are for gray scale and color images [11], where the gray level or color value of a selected group of pixels is changed by a small amount without causing visually noticeable artifacts. These techniques cannot be directly applied to binary document images where the pixels have either a 0 or a 1 value. Arbitrarily changing pixels on a binary image causes very noticeable artifacts (see Figure 1 for an example).

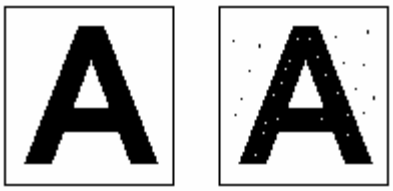

Figure 1. Effect of arbitrarily changing pixel values on a binary image

A different class of embedding techniques must therefore be developed. These would have important applications in a wide variety of document images that are represented as binary foreground and background; for example, bank checks, financial instruments, legal documents, driver licenses, birth certificates, digital books, engineering maps, architectural drawings, road maps, and so forth. Until recently, there has been little work on watermarking and data hiding techniques for binary document images. In the remaining portion of this chapter we describe some general principles and techniques for document image watermarking and data hiding. Our aim is to give the reader a better understanding of the basic principles, inherent trade-offs, strengths, and weaknesses of document image watermarking and data hiding techniques that have been developed in recent years.

3. Data Hiding Techniques for Documents Images

Watermarking and data hiding techniques for binary document images can be classified according to one of the following embedding methods: text line, word, or character shifting, fixed partitioning of the image into blocks, boundary modifications, modification of character features, modification of run-length patterns, and modifications of half-tone images. In the rest of this section we describe representative techniques for each of these methods.

1. Text Line, Word or Character Shifting:

One class of robust embedding methods shifts a text line, a group of words, or a group of characters by a small amount to embed data. They are applicable to documents with formatted text. S. Low and co-authors have published a series of papers on document watermarking based on line and word shifting [2]. These methods are applicable to documents that contain

paragraphs of printed text. Data is embedded in text documents by shifting lines and words spacing by a small amount (1/150 inch.) For instance, a text line can be moved up to encode a '1' or down to encode a '0,' a word can be moved left to encode a '1' or right to encode '0'. The techniques are robust to printing, photocopying, and scanning. In the decoding process, distortions and noise introduced by printing, photocopying and scanning are corrected and removed as much as possible. Detection is by use of maximum-likelihood detectors. In the system they implemented, line shifts are detected by the change in the distance of the marked line and two control lines — the lines immediately above and below the marked line. In computing the distance between two lines, the estimated centroids of the horizontal profiles (projections) of the two lines are used as reference points. Vertical profiles (projections) of words are used for detecting word shifts. The block of words to be marked (shifted) is situated between two control blocks of words. Shifting is detected by computing the correlation between the received profile and the uncorrupted marked profile. The line shifting approach has low embedding capacity but the embedded data are robust to severe distortions introduced by processes such as printing, photocopying, scanning, and facsimile transmission. The word shifting approach has better data embedding capacity but reduced robustness to printing, photocopying and scanning. It attempts to combine the unobtrusiveness of spatial domain techniques and the good detection performance of frequency domain techniques. Marking is performed according to the line and word shifting method described above. The frequency watermark X is then computed as the largest N values of the absolute differences in the transforms of the original document and the marked document. In the detection process, the transform of the corrupted document is first computed. The corrupted frequency watermark X* is then computed as the largest N values of the absolute differences in the transform of the corrupted document and the original document. The detection of watermark is by computing a similarity between X and X*. This method assumes that the transform of the original document and the frequency watermark X computed from the original document and the marked document (before corruption) is available during the detection process. It is shown that the height of a bounding box enclosing a group of words can be used to embed data. The height of the bounding box is increased by either shifting certain words or characters upward, or by adding pixels to end lines of characters with ascenders or descenders. The method was proposed to increase the data embedding capacity over the line and/ or word shifting methods described above. Experimental results show that bounding box expansions as small as 1/300 inch can be reliably detected after several iterations of photocopying[1]. For each mark, one or more adjacent words on an encodable text line are selected for displacement according to a selection criterion. The words immediately before and after the shifted word(s), and a block of words on the text line immediately above or below the shifted word(s), remain unchanged and are used as "reference heights" in the decoding process. The box height is measured by computing a local horizontal projection profile for the bounding box. This method is very sensitive to baseline skewing. A small rotation of the text page can cause distortions in bounding box height, even after de-skewing corrections. Proper methods to deal with skewing require further research. Character spacing is used as the basic mechanism to hide data. A line of text is first divided into blocks of characters. A data bit is then embedded by adjusting the widths of the spaces between the characters within a block, according to a predefined rule [5]. This method has advantage over the word spacing method above in that it can be applied to written languages that do not have spaces with sufficiently large width for word boundaries. The method has embedding capacity comparable to that of the word shifting method. Embedded data are detected by matching character spacing patterns corresponding to data bits '0' or '1'. Experiments show that the method can withstand document duplications. However, improvement is needed for the method to be robust against severe document degradations. This could be done by increasing the block size for embedding data bits, but this also decreases the data embedding capacity [4].

2. Fixed Partitioning of Images:

One class of embedding methods partitions an image into fixed blocks of size m x n, and computes some pixel statistics or invariants from the blocks for embedding data. They can be applied to binary document images in general; for example, documents with formatted text or engineering drawings. In Wu et al. (2000), the input binary image is divided into 3x3 (or larger) blocks. The flipping priorities of pixels in a 3x3 block are then computed and those with the lowest scores can be changed to embed data. The flipping priority of a pixel is indicative of the estimated visual distortion that would be caused by flipping the value of a pixel from 0 to 1 or from 1 to 0. It is computed by considering the change in smoothness and connectivity in a 3x3 window centered at the pixel. Smoothness is measured by the horizontal, vertical, and diagonal transitions, and connectivity is measured by the number of black and white clusters in the 3x3 window. Data are embedded in a block by modifying the total number of black pixels to be either odd or even, representing data bits 1 and 0, respectively. Shuffling is used to equalize the uneven embedding capacity over the image. It is done by random permutation of all pixels in the image after identifying the flappable pixels. An input binary image is divided into blocks of 8x8 pixels. The numbers of black and white pixels in each block are then altered to embed data bits 1 and 0. A data bit 1 is embedded if the percentage of white

pixels is greater than a given threshold, and a data bit 0 is embedded if the percentage of white pixels is less than another threshold. A group of contiguous or distributed blocks is modified by switching white pixels to black or vice versa until such thresholds are reached. For ordinary binary images, modifications are carried out at the boundary of black and white pixels, by reversing the bits that have the most neighbours with the opposite pixel value. For dithered images, modifications are distributed throughout the whole block by reversing bits that have the most neighbours with the same pixel value. This method has some robustness against noise if the difference between the thresholds for data bits 1 and 0 is sufficiently large, but this also decreases the quality of the marked document [8]. A data hiding scheme using a secret key matrix K and a weight matrix W is used to protect the hidden data in a host binary image. A host image F is first divided into blocks of size mxn. For each block Fi, data bits b1,b….br are embedded

$$SUM((F_i \oplus K) \otimes W) \equiv b_1 b_2 ... b_r \pmod{2^r},$$

where $\oplus$ represents the bit-wise exclusive OR operation, $\otimes$ represents pair-wise multiplication, and *SUM* is the sum of all elements in a matrix. Embedded data can be easily extracted by computing:

$$SUM((F_i \oplus K) \otimes W) \pmod{2^r}$$

The scheme can hide as many as $\lfloor \log_2(mn+1) \rfloor$ bits of data in each image block by changing at most two bits in the image block. It provides high security, as long as the block size (m x n) is reasonably large. In a 256x256 test image divided into blocks of size 4x4, 16,384 bits of information were embedded. This method does not provide any measure to ensure good visual quality in the marked document [13]. An enhancement was made to the method by imposing the constraint that every bit that is to be modified in a block is adjacent to another bit that has the opposite value. This improves the visual quality of the marked image by making the inserted bits less visible, at the expense of sacrificing some data hiding capacity. The new scheme can hide up to $\lfloor \log_2(mn+1) \rfloor - 1$ bits of data in an m x n image by changing at most two bits in the image block [19].

3. Boundary Modifications:

The data are embedded in the eight-connected boundary of a character. A fixed set of pairs of five-pixel long boundary patterns were used for embedding data. One of the patterns in a pair requires deletion of the center foreground pixel, whereas the other requires the addition of a foreground pixel. A unique property of the proposed method is that the two patterns in each pair are dual of each other — changing the pixel value of one pattern at the center position would result in the other. This property allows easy detection of the embedded data without referring to the original document, and without using any special enforcing techniques for detecting embedded data. Experimental results showed that the method is capable of embedding about 5.69 bits of data per character (or connected component) in a full page of text digitized at 300 dpi. The method can be applied to general document images with connected components; for example, text documents or engineering drawings [16].

4. Modifications of Character Features:

This class of techniques extracts local features from text characters. Alterations are then made to the character features to embed data. Text areas in an image are identified first by connected component analysis, and are grouped according to spatial closeness. Each group has a bounding box that is divided into four partitions [14]. The four partitions are divided into two sets. The average width of the horizontal strokes of characters is computed as feature. To compute average stroke width, vertical black runs with lengths less than a threshold are selected and averaged. Two operations — "make fat" and "make thin" — are defined by increasing and decreasing the lengths of the selected runs, respectively. To embed a "1" bit, the "make fat" operation is applied to partitions belonging to set 1, and the "make thin" operation is applied to partitions belongs to set 2. The opposite operations are used to embed "0" bit. In the detection process, detection of text line bounding boxes, partitioning, and grouping are performed. The stroke width features are extracted from the partitions, and added up for each set. If the difference of the sum totals is larger than a positive threshold, the detection process outputs 1. If the difference is less than a negative threshold, it outputs 0. This method could survive the distortions caused by print-and-scan (redigitization) processes. The method's robustness to photocopying needs to be furthered investigated. A scheme is presented to embed secret messages in the scanned gray scale image of a document. Small sub-characterized regions that consist of pixels that meet criteria of text-character parts are identified first, and the lightness of these regions are modulated to embed data. The method employs two scans of the document — a low resolution scan and a high resolution scan. The low-resolution scan is used to identify the various components of the document and establish a coordinate system based on the paragraphs, lines and words found in the document [20]. A list of sites for embedding data is selected from the low resolution scanned image. Two site selection methods were presented in the paper. In the first method, a text paragraph is partitioned into grids of 3x3 pixels. Grid cells that contain predominately text-type pixels are selected. In the second method, characters with long strokes are identified. Sites are selected at locations along the stroke. The second scan is a

full-resolution scan that is used to generate the document copy. The pixels from the site lists generated in the low-resolution scan are identified and modulated by the data bits to be embedded. Two or more candidate sites are required for embedding each bit. For example, if the difference between the average luminance of the pixels belonging to the current site and the next one is positive, the bit is a 1; else, the bit is a 0. For robustness, the data to be embedded are first coded using an error correcting code. The resulting bits are then scrambled and dispersed uniformly across the document page. For data retrieval, the average luminance for the pixels in each site is computed and the data are retrieved according to the embedding scheme and the input site list. This method was claimed to be robust against printing and scanning. However, this method requires that the scanned grayscale image of a document be available. The data hiding capacity of this method depends on the number of sites available on the image, and in some cases, there might not be enough sites available to embed large messages.

5. Modification of Run-Length:

A method was proposed to embed data in the run-lengths of facsimile images. A facsimile document contains 1,728 pixels in each horizontal scan line. Each run length of black (or foreground) pixels is coded using modified Huffman coding scheme according to the statistical distribution of run-lengths. In the proposed method, each run length of black pixels is shortened or lengthened by one pixel according to a sequence of signature bits. The signature bits are embedded at the boundary of the run lengths according to some pre-defined rules [5].

6. Modifications of Half-Toned Images:

Several watermarking techniques have been developed for half-tone images that can be found routinely in printed matters such as books, magazines, newspapers, printer outputs, and so forth. This class of methods can only be used for half-tone images, and are not suitable for other types of document images. The methods to embed data during the half-toning process. This requires the original gray scale image. The methods describe to embed data directly into the half-tone images after they have been generated. The original grayscale image is therefore not required. A sequence of two different dither matrices (instead of one) was used in the half-toning process to encode the watermark information. The order in which the two matrices are applied is the binary representation of the watermark [16]. Three methods were proposed to embedded data at pseudo-random locations in half-tone images without knowledge of the original multi-tone image and the half-toning method. The three methods, named DHST, DHPT, and DHSPT, use one half-tone pixel to store one data bit. In DHST, N data bits are hidden at N pseudo-random locations by forced toggling. That is, when the original half-tone pixel at the pseudo-random locations differs from the desired value, it is forced to toggle. This method results in undesirable clusters of white or black pixels. In the detection process, the data are simply read from the N pseudo-random locations. In DHPT, pair of white and black pixels (instead of one in DHST) is chosen to toggle at the pseudo-random locations. This improves over DHST by preserving local intensity and reducing the number of undesirable clusters of white or black pixels. DHSPT improves upon DHPT by choosing airs of white and black pixels that are maximally connected with neighboring pixels before toggling. The chosen maximally connected pixels will become least connected after toggling and the resulting clusters will be smaller, thus improving visual quality. An algorithm called intensity selection (IS) is proposed to select the best location, out of a set of candidate locations, for the application of the DHST, DHPT and DHSPT algorithms. By doing so, significant improvement in visual quality can be obtained in the output images without sacrificing data hiding capacity. In general, the algorithm chooses pixel locations that are either very bright or very dark. It represents a data bit as the parity of the sum of the half-tone pixels at M pseudo-random locations and selects the best out of the M possible locations. This algorithm, however, requires the original grayscale image or computation of the inverse-half-toned image [12]. Two data hiding techniques for digital half-tone images were described: modified ordered dithering and modified multiscale error diffusion. In the first method, one of the 16 neighboring pixels used in the dithering process is replaced in an ordered or pre-programmed manner. The method was claimed to be similar to replacing the insignificant one or two bits of a grayscale image, and is capable of embedding 4,096 bits in an image of size 256 x 256 pixels. The second method is a modification of the multi-scale error diffusion (MSED) algorithm for half-toning as proposed which alters the binarization sequence of the error diffusion process based on the global and local properties of intensity in the input image. The modified algorithm uses fewer floors (e.g., three or four) in the image pyramid and displays the binarization sequence in a more uniform and progressive way. After 50% of binarization is completed, the other 50% is used for encoding the hidden data. It is feasible that edge information can be retained with this method. A joint half toning and watermarking approach that combines optimization based halftoning with a spread spectrum robust

watermark. The method uses a joint metric to account for the distortion between a continuous tone and a halftone (FWMSE), as well as a watermark detectability criterion (correlation). The direct binary search method is used for searching a halftone that minimizes the metric. This method is obviously extendable in that other distortion metric and/ or watermarking algorithms can be used [13].

4. CONCLUSIONS

We have presented an overview and summary of recent developments in binary document image watermarking and data hiding research. Although there has been little work done on this topic until recent years, we are seeing a growing number of papers proposing a variety of new techniques and ideas. Research on binary document watermarking and data hiding is still not as mature as for color and grayscale images. More effort is needed to address this important topic. Future research should aim at finding methods that offer robustness to printing, scanning, and copying, yet provide good data embedding capacity. Quantitative methods should also be developed to evaluate the quality of marked images. The steganographic capability of different techniques needs to be investigated and techniques that can be used in covert communication applications need to be developed.

5. REFERENCES

[1]Allebach J.P., Flohr T.J., Hilgenberg D.P. & Atkins, C.B. (1994, May). "Model-based half toning via direct binary search". Proceedings of IS&T's 47th Annual Conference, (pp. 476-482), Rochester, NY.

[2]Amamo, T. & Misaki D. (1999). "Feature calibration method for watermarking of document images". Proceedings of 5th Int'l Conference on Document Analysis and Recognition, (pp. 91-94), Bangalore, India.

[3]Baharav Z. & Shaked D. (1999, January). "Watermarking of dither half-toned images". Proc. of SPIE Security and Watermarking of Multimedia Contents, 1,307-313.

[4]Bhattacharjya A.K. & Ancin H. (1999). "Data embedding in text for a copier system". Proceedings of IEEE International Conference on Image Processing, 2, 245-249.

[5]Brassil J. & O'Gorman L. (1996, May). "Watermarking document images with bounding box expansion". Proceedings of 1st Int'l Workshop on Information Hiding, (pp. 227-235). Newton Institute, Cambridge, UK.

[6]Chotikakamthorn N. (1999). "Document image data hiding techniques using character spacing width sequence coding". Proc. IEEE Intl. Conf. Image Processing, Japan.

[7]Cox I., Kilian J., Leighton T. & Shamoon T. (1996, May/June). "Secure spread spectrum watermarking for multimedia". In R. Anderson (Ed.), Proc. First Int. Workshop Information Hiding (pp. 183-206). Cambridge, UK

[8]Craver S., Memon N., Yeo B. & Yeung M. (1998, May). "Resolving rightful ownership with invisible watermarking techniques: Limitations, attacks, and implications". IEEE Journal on Selected Areas in Communications, 16(4), 573-586.

[9]Digimarc Corporation. http://www.digimarc.com.

[10]Foley J.D., Van Dam A., Feiner S.K. & Hughes J.F. (1990). "Computer graphics: Principles and practice (2nd ed.). Addison-Wesley".

[11]Fu M.S. & Au O.C. (2000a, January). "Data hiding for halftone images". Proc of SPIE Conf. On "Security and Watermarking of Multimedia Contents II", 3971, 228-236.

[12]Fu M.S., & Au O.C. (2000b, June 5-9). "Data hiding by smart pair toggling for halftone images". Proc. of IEEE Int'l Conf. Acoustics, Speech, and Signal Processing, 4, (pp. 2318-2321).

[13]Fu M.S. & Au O.C. (2001). "Improved halftone image data hiding with intensity selection". Proc. IEEE International Symposium on Circuits and Systems, 5, 243-246.

[14]Holliman M. & Memon N. (2000, March). "Counterfeiting attacks and block wise independent watermarking techniques". IEEE Transactions on Image Processing, 9(3), 432-441.

[15]Kacker D. & Allebach, J.P. (2003, April). "Joint half toning and watermarking". IEEE Trans. Signal Processing, 51, 1054-1068.

[16]Katsavounidis I. & Jay Kuo C.C. (1997, March). "A multiscale error diffusion technique for digital half–toning". IEEE Trans. on Image Processing, 6(3), 483-490.

[17]Knox K.T. "Digital watermarking using stochastic screen patterns", United States Patent Number 5,734,752.

[18]Koch E. & Zhao J. (1995, August). "Embedding robust labels into images for copyright protection". Proc. International Congress on Intellectual Property Rights for Specialized Information, Knowledge & New Technologies, Vienna.

[18]Liu Y., Mant J., Wong E. & Low, S.H. (1999, January). "Marking and detection of text documents using transform-domain techniques". Proc. SPIE Conf. on Security and Watermarking of Multimedia Contents, (pp. 317-328),San Jose, CA.

[19]Low S.H., Lapone A.M. & Maxmchuk N.F. (1995, November 13-17). "Document identification to discourage illicit copying". IEEE GlobeCom 95, Singapore.

[20]Low S.H. & Maxmchuk N.F. (1998, May). "Performance comparison of two text marking methods". IEEE Journal on Selected Areas in Communications, 16(4).